The Great Compression vs the Clean Transition: How the Great Recession held back clean innovation*

Ralf Martin† Dennis Verhoeven‡

May, 2024

Abstract

We examine the distribution of the private value of clean and dirty innovation using new methods based on patent data. We document that the value of clean innovations is higher and more dispersed. We find an overall decline in the variability of private values and returns in the wake of the Great Recession. This is consistent with the idea that financial restrictions have made investors and innovators more risk averse. Because clean and dirty innovations show different exposures to such risk aversion, the recession could have contributed to the decline of clean relative to dirty innovation. We develop a method to quantify counterfactual clean and dirty innovation that would have prevailed if the distribution of private values (or returns) had stayed fixed. The results suggest that shying away from risky R&D in the wake of the Great Recession has considerably depressed the relative share of clean innovation. More broadly, our results suggest that financial constraints in the wake of crises may be an important barrier on the path to the clean equilibrium.

Keywords: Innovation, R&D, Climate Change, Clean Transition

*We thank Robert Baluja for excellent research assistance. Verhoeven gratefully acknowledges funding from the Research Foundation Flanders (FWO grant 1241520N) and the H2020 Framework Programme (Marie Skłodowska-Curie Action DIV_INV grant 844254)

†International Finance Corporation, Imperial College, LSE and CEPR, r.martin@imperial.ac.uk

‡Bocconi University, ICRIOS; KU Leuven, Department of Management, Strategy & Innovation (MSI); London School of Economics, Centre for Economic Performance (CEP); Research Foundation Flanders (FWO) — contact: dennis.verhoeven@kuleuven.be.

1 Introduction

Technological innovation in the production of clean energy is a key element in the transition to a net-zero economy. Given important environmental and innovation externalities, the private sector is unlikely to invest adequate resources in producing such innovations, and this has generated substantial interest in policies that increase clean energy innovation. Recent studies have documented a substantial drop in Clean (patented) innovation starting in 2010, with negative annual growth rates of 5.5% (Popp et al., 2022; Probst et al., 2021; Acemoglu et al., 2019). Much less is known about the *relative* growth path of Clean as opposed to Dirty innovation. We view this as important in the light of models that show path-dependency in the direction of technological change based on extent of the accumulated stock of knowledge (Acemoglu et al., 2012, 2016). When the Clean knowledge stock grows faster – either because of market conditions or government intervention – than the Dirty one, the equilibrium naturally shifts to the further development of sustainable technology.

This paper explores what the evolution of private value of Clean and Dirty innovation can teach us about this relative growth path. Measures of private value allow us to attach weights to individual innovations in order to separate important from frivolous inventions. More importantly, looking at the evolution of the dispersion in private value of innovation sheds light on the risk appetite of firms pursuing R&D projects. This allows us to explore the role of financial constraints in the wake of the Great Recession in shaping the relative growth of Clean innovation.

We rely on the methodology proposed by Kogan et al. (2017) (henceforth KPSS) to measure private value. This market-based approach values innovation by measuring market-adjusted returns to a firm's stock around the day a patent grant is announced. This event study approach is attractive in that it recovers how the market responds to the news of a successful patent application. However, it has two important drawbacks when applied to our purpose. First, it suffers from the standard bias in event study analysis: the change in the stock price only recovers the full value of new information if that information is a complete surprise. If the market anticipates that a firm could be granted a patent, then the stock price will have already internalized some of the value of the potential patent before the grant announcement is made. Adjusting for this requires estimating the proba-

bility with which the market expects a patent grant. KPSS assume a fixed probability of 0.56, which is equal to the observed frequency of successful patent applications between 1991-2001. However, this probability may vastly differ based on the underlying technology. Some sectors may have a harder time obtaining patents than do others. If the market incorporates these differences in probabilities, the estimated value of patented innovations could be off by a considerable margin. We address this problem by adjusting private values based on field-specific grant rates. A second drawback is that the market response to a patent grant does not account for the cost of an innovation project, as it is sunk at the time of grant. At the same time, what matters for our analysis of risk appetite is not the value of an innovation as such but the return to R&D investments. We use a new approach proposed by Guillard et al. (2021)) to recover an estimate of the cost of R&D investments by fitting a simple model of innovation to the distribution of private values.

We use our measures to show descriptive patterns of the evolution of Clean and Dirty innovation. We document a rise followed by a drop in the number and the total value of both Clean and Dirty innovation. As this pattern is more pronounced for Clean, its relative share follows a similar rise-then-fall trend. We also find that the average private value and return rate of Clean innovation is slightly higher and more dispersed than that of Dirty (and overall) innovation. Over time - particularly after the onset of the Great recession - there is a considerable decline in both the value and returns dispersion for clean and dirty innovation. These descriptive patterns are consistent with investors seeking lower risk in the wake of the Great Recession. Given different exposures of Clean and Dirty innovations to the outer ends of the value and return distributions, such increase in risk-aversion could have impacted the relative growth path.

We explore this further by deriving counterfactual innovation results by fixing the pre-recession value or return distribution of overall innovation. Assuming that changes in the overall dispersion reflect changes in risk appetite due to the recession, this approach allows us to estimate the counterfactual relative growth path. Results from this exercise suggest a considerable impact of such financial restrictions on the relative growth path of Clean innovation. Indeed, even our most conservative scenario would suggest that the 2014 counterfactual share of Clean relative to Dirty innovation would be equal to its 2011 actual peak.

While we think the financial constraints explanation buys us traction in explaining changes to relative Clean growth, our analyses do not rule a number of other explanations. First, changes in fuel prices or the emergence of new fossil fuel technologies (Acemoglu et al., 2019) may have increased the attraction of 'dirty' technologies. Second, decreasing real wages in the wake of the financial crisis may have undermined the willingness to pay for clean energy products. Third, the wave of clean energy innovation preceding the recent slump (i.e. from the 1990 to around 2011) might have been so successful that it exhausted easy innovation options and 'clean' ideas are now harder to find (Popp et al., 2022). While much further analysis is needed, we think our results justify further exploration of how crisis-induced risk appetite shifts may shape the relative growth path of Clean innovation.

2 Methods and Data

2.1 Data

We collect patent information from the PATSTAT Global database.[1] This database is curated by the European Patent Office and collects information from patent offices worldwide. We focus on inventions for which a patent was filed at the United States Patent and Trademark Office (USPTO). Inventions are timestamped based on the first patent filing related to an invention, even if that first filing occurred at a jurisdiction outside the USPTO. This is appropriate because the earliest filing date is the best approximation of when the actual invention occurred.

Our classification of inventions into Clean, Dirty and Other categories relies on the Cooperative Patent Classification (CPC). The CPC is a hierarchical classification system jointly developed by the EPO and the USPTO to help patent examiners identify relevant prior patents against which the novelty and inventive step of new applications are judged. For each patent application, the examiner assigns one or multiple CPC classes. At its most detailed level, it contains around 250,000 classes and therefore serves as an invaluable taxonomy of the technological landscape.[2] To assign patents to various Clean technologies, we use a mapping developed at the UK government Bureau of Energy and Industrial Strategy

[1] Spring 2018 version

[2] The CPC has a particular advantage when it comes to identifying clean inventions because it labels them using the class: '*Y02: Technologies or applications for mitigation or adaptation against climate change*'.

(BEIS). This mapping is the result of an extensive effort to identify CPC classes related to various clean innovation sub-fields. It has been cross-validated against mappings in the academic literature and by engineering experts. To assign patents to 'Dirty' technologies, we use the mapping of CPCs developed in Dechezlepretre et al. (2020). When none of a patent's classes is assigned to the labels 'Dirty' or 'Clean', we classify it as 'Other'.[3]

2.2 Event study approach to measure private value

We rely on a measure for the private value of patented inventions developed in KPSS. It relies on an event study approach that backs out the private value from abnormal stock returns during the 3-day window around the grant of a patent. Abnormal stock returns are defined as the return on the patenting firm's share relative to the return on the market portfolio. As such, abnormal stock returns filter out all market-wide shocks to stock prices. To filter out firm-specific shocks unrelated to the grant of the patent, they rely on assumptions about the distribution of returns that can be ascribed to patent grants and returns because of unrelated events.[4] Estimating the parameters of these distributions then allows us to formulate an expression of the expected returns that can be attributed to the grant of the patent.

The event study approach is attractive because it recovers how the market responded to the information that an invention made by the firm resulted in a patent right. It provides a measure for the future profits the market expects the innovation to generate at the time a patent is granted. However, the patent-induced returns suffer from the standard bias in event study analyses in that it only recovers the value of information to the extent it is a surprise to the market. To work around the issue that the value of the patented invention – should it be granted – may be common knowledge to the market, KPSS multiply the returns by $(1-\pi)^{-1}$, where $\pi = 0.56$ is an estimate of the probability with which the market believes the patent will be granted. It corresponds to the overall grant rate of patents at the USPTO between 1991 and 2001. To convert the resulting estimated patent-specific return to a dollar-value, it is multiplied by the market capitalization of the patenting firm.

[3] As patents are often assigned multiple CPC codes, they may be labeled as both 'Dirty' and 'Clean'. In such cases, we choose the label that occurred most often, and a random one in case of a tie.

[4] Specifically, they assume returns related to a patent grant are positive and follow a normal distribution truncated at zero, and that idiosyncratic returns are normally distributed around zero.

We propose an adjustment of this measure that takes into account how grant probabilities vary across technological fields. To make this adjustment, we calculate the USPTO grant rate across the categories defined for patents filed between 2005 and 2009.[5] Restricting to the first half of the sample gives each patent at least 5 years to be granted. This results in alternative values for π that vary across fields, and is used to adjust private values by multiplying it with $\frac{(1-\pi_{new})^{-1}}{(1-\pi_{old})^{-1}}$. Table 1 shows the grant rate by field and the multiplier applied to private values obtained from KPSS. Appendix A provides further details on the calculation of (adjusted) private value.

Table 1: Adjusting for grant probability

Field	Grant probability	Nr. innovations	Multiplier
Solar	0.46	3616	0.82
Heating and Cooling	0.53	6488	0.93
Biomass & Bioenergy	0.59	1369	1.07
Tidal Stream	0.6	189	1.09
Offshore Wind	0.63	2750	1.19
Other	0.64	1262482	1.21
Industry	0.64	8016	1.22
CCUS	0.65	346	1.25
Smart Systems	0.65	2464	1.27
Hydrogen	0.67	1975	1.33
Building Fabric	0.68	12496	1.37
Dirty Electricity	0.71	18166	1.51
Nuclear	0.77	3827	1.88
Dirty Transport	0.78	7774	2.03

Notes: Deriving the multiplier that adjusts private values for grant probability. Columns 1 and 2 show the Clean, Dirty and Other categories and the probability any USPTO patent filed for between 2005 and 2009 was granted by 2014. Third column shows the number of applications the grant probability was calculated on. Fourth column shows the multiplier applied to private value estimates of innovations in each category. It is equal to $\frac{(1-\pi_{new})^{-1}}{(1-\pi_{old})^{-1}}$, where π_{new} is equal to the grant probability in column 2 and $\pi_{old} = 0.56$ (the grant probability assumed in KPSS)

Table 1 shows that there is substantial variation in the likelihood that a filed patent is granted. In established industries such as dirty electricity, nuclear, and dirty transport, over 70% of filed patents are granted within five years. On the other hand, patents filed in solar or heating and cooling have a 46% and 53% chance of being granted within five years. These differences matter substantially for the implied value of patents. For instance, using the estimate of the likelihood of a patent being granted from KPSS under-states the value of dirty transport patents and needs to be inflated by a factor of 2.03. This is because more of the value of these patents is "priced in" to stock prices before patent

[5]To calculate the grant rate we apply a window of 5 years from the first filing at the USPTO. This means that patents granted over 5 years after application are considered non-granted. This approach ensures comparability of earlier and later patents in the sample. From all granted patents in our sample, 96.3% are granted within 5 years. Not applying a 5-window does not materially change the results.

announcement since it is very likely that the patent will be approved. On the other hand, KPSS somewhat over-states the value of solar patents. The large spread in the multipliers in Table 1 suggests that correctly accounting for the probability of a patent being granted could change our understanding of the recent drop in green patenting.

2.3 Inferring R&D cost

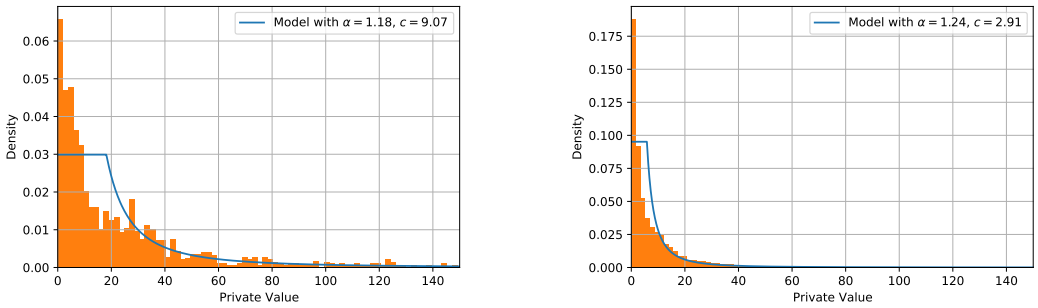
The (adjusted) private value of a patent captures the market's estimate of the profits that will accrue to the firm from its resulting monopoly right on the day of grant. On that day, the costs of R&D are sunk, so the private value includes the value created (through future turnover) but not the innovation costs, and so not the profits related to the invention. For our analyses on the impact of the Great Recession on risk appetite, we are interested in firms' incentives provided by the returns to innovation. We therefore need to compare the value of an innovation with its R&D investment. Unfortunately, there is no data set that would give us comprehensive information on this for all - or even a large fraction of - the innovations we can observe in our patent data.

To circumvent this problem, we apply the approach developed in Guillard et al. (2021). They introduce a simple model (see Appendix B for further details) of the innovation process that allows us to back out the costs in a technological field from the observed private value distribution. In their model, innovators receive ideas randomly according to a left-skewed Pareto distribution with a parameter α that defines its skew. For ideas to turn into an innovation, innovators have to invest a fixed cost c. Before investing, the innovator observes the quality of the idea. Investing results in an uncertain profit that depends on the idea quality. The need to invest a fixed cost means that the inventor will only pursue ideas of sufficient quality. These assumptions produce a kinked distribution of the private value of ideas ex post (those that we observe in the data). To infer the unobserved cost, they use the kink in the distribution of observed private values. Specifically, they estimate the cost (and shape parameter α) by fitting the modelled distribution to the observed distribution of private values for each of over 600 technology classes (CPC subclasses) for

each year between 2005 and 2014. Figure 1 shows the result of this procedure for two of the most prevalent subclasses. For class *A61K* (Pharmaceuticals) there is more probability mass on higher values as compared to class *G06F* (Data processing). The fitted model distribution reflects this pattern and assigns a higher cost to this field.[6]

Figure 1: Estimating costs from PV distributions

(a) A61K **(b)** G06F



Notes: Comparison of actual and modeled private value distributions for two prevalent IPC subclasses (*A61K: preparations for medical, dental, or toilet purposes* and *G06F: Electric digital data processing*). Histogram plots actual private value distribution in the class, blue line shows the modeled density.

3 Trends in (the value of) Clean and Dirty innovation

3.1 Clean innovation drop

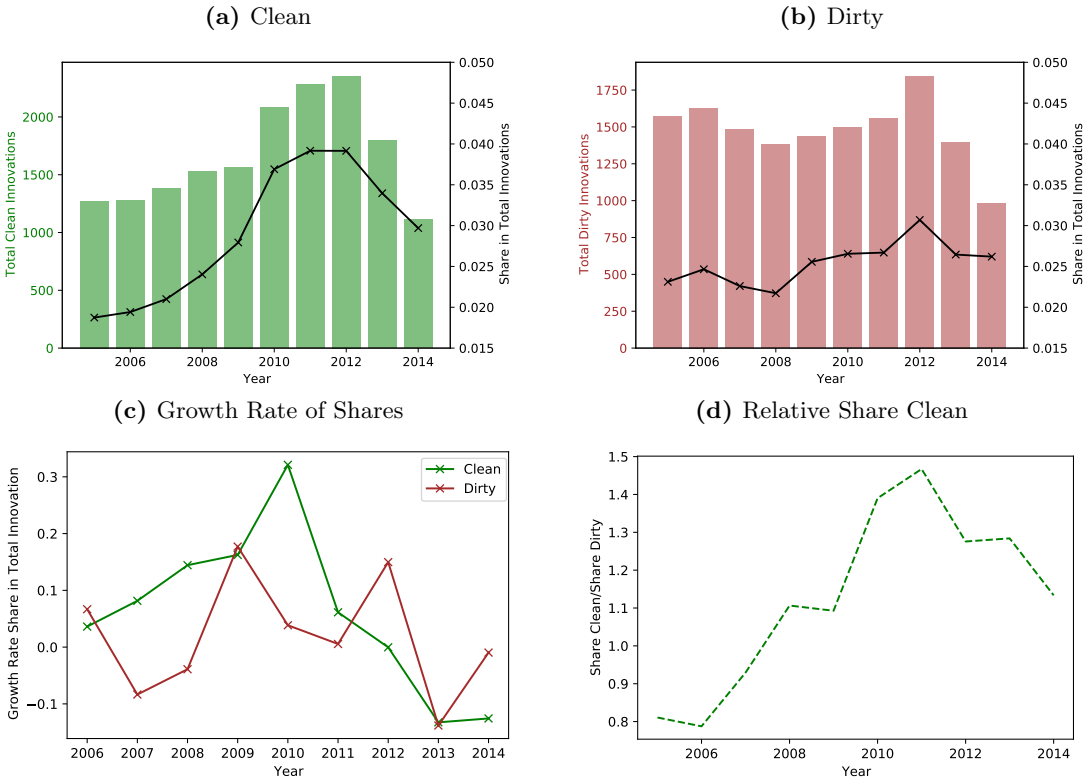
Figure 2 illustrates recent trends in clean innovation. We restrict our focus to innovations from listed US firms, but Appendix C (Figure 11) shows trends for all innovations with a USPTO application. We see a steep increase in both the absolute number and share of clean innovations between 2005 and 2010. The share of Clean innovation nearly doubled during this period, from 1.9% to about 3.9%. Growth levelled off between 2010 and 2012, and strongly dropped in 2013 and 2014, to a share of 3%. The average yearly growth rate between 2005 and 2011 was 13.4%, while that between 2012 and 2014 was -4.8% (which is similar to the drop observed in previous work). For the share of Dirty innovations we also see an upward trend up until 2012, and a drop after. These trends, however, are not nearly as pronounced as for Clean. Especially the increase before 2012 was much less pronounced (with an increase in shares from about 2.3% to 3.1%). The drop in shares the last two

[6]The estimated cost is equal to half the value at which the horizontal straight blue line 'kinks' into the curve.

years of our sample was also less pronounced as compared to Clean, from 3.1% to 2.6%. To investigate the *relative* trends, panel (d) shows the relative share of Clean innovation (dividing Clean by Dirty shares). Up until 2009, there was a bit less than one Clean innovation for each Dirty one. The relative share of Clean grew to about 1.5 to 1 by 2010, but this growth has not been sustained into the 2010s. Between 2011 and 2014, there have been about 1.2 Clean innovations for each Dirty one. This trend is concerning in the light of the transition to a clean economic equilibrium. To achieve such equilibrium, we need that clean knowledge accumulation exceeds dirty knowledge accumulation (Acemoglu et at 2012).

Figure 3 repeats the exercise when weighting innovations for their private value (adjusted for grant probability). The share of total value generated by Clean innovations follows a similar pattern to that of counts of innovations, albeit somewhat less pronounced. The share was rather stable between 2005 and 2008 (while it was increasing when looking at innovation counts), but increased to about 4% in 2010 (similar to the share in counts), after which it started decreasing to about 3.5%. This drop is less pronounced than for innovation counts (where it dropped to about 3%). For Dirty innovation, weighting for private value results in a gradually decreasing trend, but the 2012 peak and subsequent drop remains visible. Because of this drop, the relative shares paint a slightly more optimistic picture with respect to the clean transition, with relative shares picking up between 2012 (but not reaching the levels of the peak in 2010).

Figure 2: Trends Clean and Dirty innovations



Notes: Trends of Clean and Dirty innovation based on the sample of all inventions for which at least one patent was filed at the USPTO, and for which the applicant was a US listed firm. Panels (a) and (b) show absolute trends in Clean and Dirty innovation respectively. The bars (left-hand side y-axis) show the number of innovations. The black lines (right-hand side y-axis) correspond to the share in all innovations (Clean, Dirty, and Other). Panel (c) shows the year-to-year growth rates of share of Clean and Dirty in all innovations. Panel (d) shows the share of Clean relative to the share of Dirty.

Figure 3: Trends Clean and Dirty total private value



Notes: Trends of Clean and Dirty total private value based on the sample of all inventions for which at least one patent was filed at the USPTO, and for which the applicant was a US listed firm. Private value calculations are adjusted for grant probability. Panels (a) and (b) show absolute trends in Clean and Dirty total private value respectively. The bars (left-hand side y-axis) show the total value in millions of CPI-adjusted 1982 US dollars. The black lines (right-hand side y-axis) correspond to the share in all innovations (Clean, Dirty, and Other). Panel (c) shows the year-to-year growth rates of share of Clean and Dirty in all innovations. Panel (d) shows the share of Clean relative to the share of Dirty.

3.2 The dispersion of values and returns

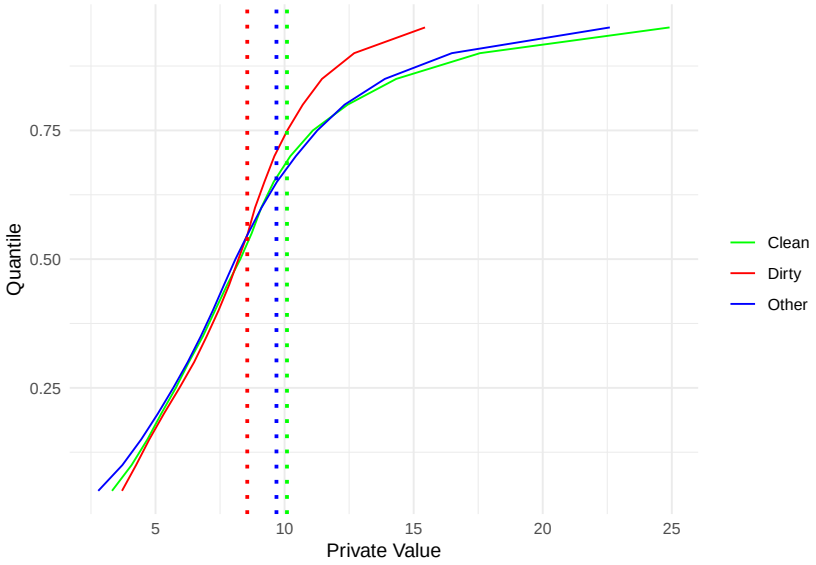
This section explores the (trends in) dispersion of private value for Clean and Dirty innovation. To get an overall view of the private value distributions, Figure 4 shows the private value for Dirty, Clean, and Other innovations for 20 quantiles of their distribution. Clean innovations are on average more valuable than both Dirty and Other innovations. Clean is also more dispersed with a greater fraction of the distribution on very low and very high values.

In figure 5 we examine dispersion trends of adjusted private values over time (Figure 12 in Appendix C repeats this figure for unadjusted private values). We normalize the values with the average private value of the category Other (all innovations not classified as Clean or Dirty) in a given year. For both Clean and Dirty, we see a strong compression of the

dispersion from 2008 onward. The value of 90th percentile for Clean reduces from about 3.5 times the value of the average Other innovation to around 2.5. For Dirty, a similar decrease can be observed, with a drop from about 3.2 to about 2. At the same time, we also see that the bottom end of the distribution moves closer to the mean for both Clean and (especially) Dirty. In Figure 6 we repeat the same exercise for innovation returns (adjusted private value over costs). For Clean, the dispersion decrease is less pronounced, while for Dirty it looks more gradual but equal in magnitude. A similar pattern is observed in Figure 13 of Appendix C, which repeats the analysis without adjusting for grant probability when calculating returns.

These figures suggest that the Clean and Dirty innovations pursued post-recession are less likely to be particularly valuable or invaluable. To more directly test this, Figure 7 plots the Clean and Dirty shares in total innovation across 5 return quintiles. We use the 2005 (adjusted) returns distribution to define 5 equal quintile bins, and assign all innovations into these bins. The figure plots the share of Clean and Dirty belonging to each returns bin. We see that the 'middle bin' (40th to 60th percentile) – plotted in red – becomes relatively more prevalent both for Clean and Dirty. The outermost bins (0-20th percentile and 80th-100th percentile) become distinctly less important. Interestingly, this wedge between 'extreme' and 'middle' bins starts at the onset of the Great Recession. This may suggest that financial constraints have reduced risk appetite in the wake of the crisis. In the following section, we explore to what extent such financial constraints may have affected the relative growth path of Clean and Dirty innovation.

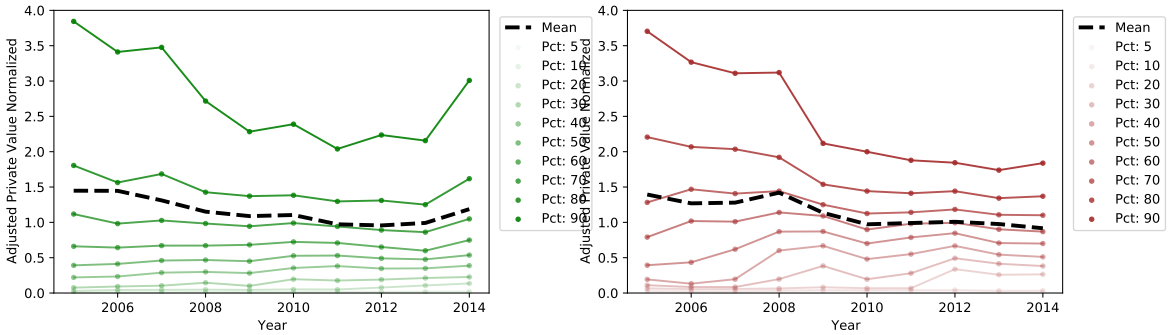
Figure 4: Private Value dispersion



Notes: Distributions of private value for Clean, Dirty and Other. Y-axis shows percentiles (in 5 percentage point steps) of each distribution. X-axis corresponds to the private value at each quantile in millions of CPI-adjusted 1982 US dollars. Vertical lines show the average value for each group.

Figure 5: Trends Adjusted Private Value dispersion

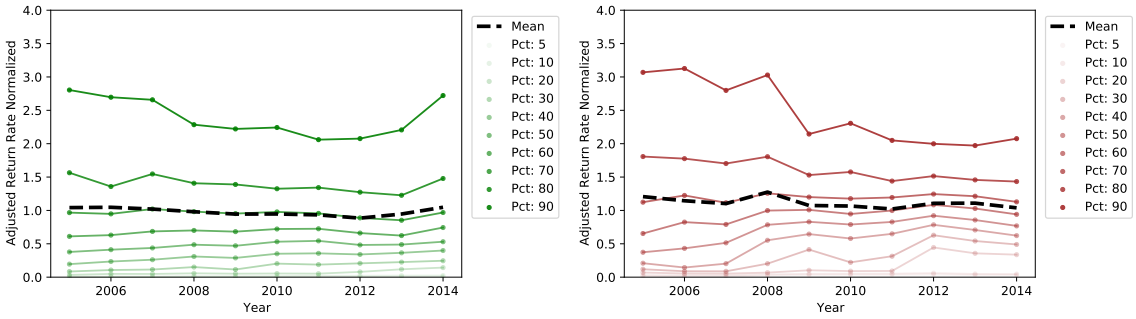
(a) Clean **(b)** Dirty



Notes: Evolution of the dispersion of private value (adjusted for grant probability in the respective category) in Clean (left) and Dirty (right). Colored lines show the evolution of each decile value (y-axis) between 2005 and 2014. Black line plots the average value for the corresponding category. Private values are normalized by the average private value for the entire sample in each year.

Figure 6: Trends Adjusted Returns dispersion

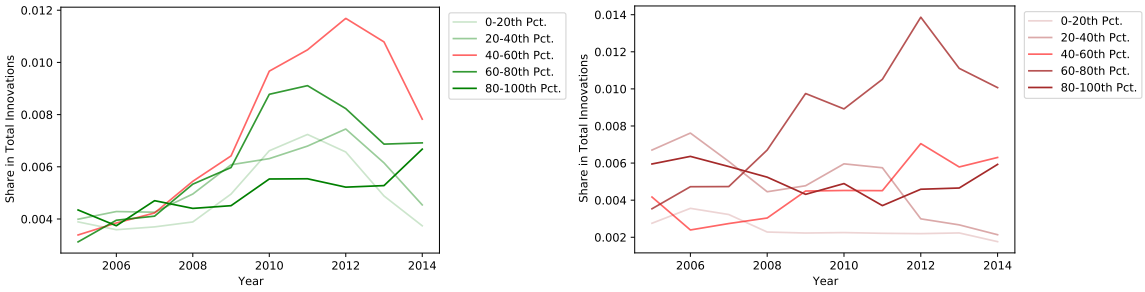
(a) Clean **(b)** Dirty



Notes: Evolution of the dispersion of returns (where private values are adjusted for grant probability in the respective category) in Clean (left) and Dirty (right). Colored lines show the evolution of each decile value (y-axis) between 2005 and 2014. Black line plots the average value for the corresponding category. Returns are normalized by the average return for the entire sample in each year.

Figure 7: Trends in the returns distribution for Clean and Dirty innovations

(a) Clean **(b)** Dirty



Notes: Shares of Clean and Dirty in all innovations across the returns distribution. The share of Clean (left) and Dirty (right) in all innovations in a given year is plotted on the y-axis. Different lines correspond to quintile bins of the overall 2005 adjusted returns distribution.

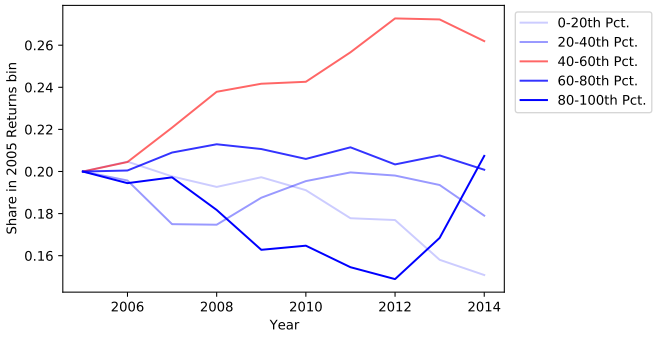
4 The impact of financial constraints

In the previous section we saw that after the Great Recession the dispersion of both Clean and Dirty innovation returns reduced. This would be consistent with the idea that financial restrictions or changes in risk preferences led innovators to avoid the more risky research projects. As we have seen, Clean and Dirty have a different exposure to the outer ends of the value and returns distribution. Hence, such a general reduction in risk appetite could have been - at least in part - responsible for the relative decline in Clean innovation.

Figure 8 checks whether a reallocation towards less risky projects happened for innovation in general. It shows a reallocation of innovation projects towards the central part of the returns distribution. For innovation in general – like for Clean and to a lesser extent Dirty innovation – growth after 2008 is primarily coming from the center of the distribution.

The Great Recession seems to have been accompanied by a Great Concentration (of the returns to R&D), which is consistent with a more risk averse attitude by investors. This raises the question of whether the relative decline of Clean innovation – having potentially less 'safe bets' – could be in part explained by a general decline in risk appetite. In the following section, we perform a simple counterfactual calculation aimed at quantifying how much of the relative Clean drop this hypothesis could explain.

Figure 8: Trends in returns distribution for all innovations



Notes: Shows the dynamics in the overall returns distribution. We define 5 quintile bins of equal size based on the 2005 adjusted returns distribution. Then we use the bin edges to assign post-2005 innovations into these 5 bins. The y-axis shows the share of innovations that belong to each of the 2005 quintile bins.

4.1 Counterfactual calculations

We make the assumption that financial constraints explain the contraction of the returns (or private value) distribution observed in the population of innovations. We then re-weight all innovations to compute counterfactual shares of clean (and dirty) innovation as follows. We take our first sample year (2005) as the baseline and ask how the composition of technology classes would have looked if the distribution of returns would have remained fixed from 2005 onward. In other words: we assume that any changes to the returns (or private value) distribution after 2005 were due to changes in the financial environment. We implement this by splitting the 2005 distribution in 100 percentile bins. For every

period $t > 2005$ we compute re-weighting factors

$$\rho_{b,t} = \frac{s_{b,2005}}{s_{b,t}}$$

where $s_{b,t}$ is the share of innovations in bin b in period t. We then compute the counterfactual shares of technology type $a \in \{Clean, Dirty\}$ as

$$s_{a,t}^{CF(x)} = \frac{\sum_{i \in a} \rho_{b(i),t}}{\sum_{i} \rho_{b(i),t}}$$

where $x \in \{PV, Return\}$; i.e. we can define the bins b on the basis of the private value or returns distribution.

In Figure 9 we report results for $s_{a,t}^{CF(PV)}$. Figure 10 we report results for $s_{a,t}^{CF(Return)}$. In both cases we see that the counterfactual innovation figures increase the clean share and depresses the dirty share. On net (see Panel d) the effect is so strong that the counterfactual clean to dirty ratio at the end of the sample is equivalent to the actual peak. However, note that in the counterfactual case also the peak is higher. In other words: financial constraints in the wake of the Great Recession have been depressing clean relative to dirty innovation ever since 2008. This would have amplified the impact of the clean drop however, there are likely other factors that have been causing the clean drop.

Figure 9: Trends re-weighted based on Private Value distribution



Notes: Trends of Clean and Dirty innovation re-weighted for the private value distribution. Private value calculations are adjusted for grant probability. Panels (a) and (b) show absolute counterfactual trends in Clean and Dirty innovation respectively. The bars (left-hand side y-axis) show the counterfactual counts. The black lines (right-hand side y-axis) correspond to the counterfactual share ($s_{a,t}^{CF(PV)}$) in all innovations (Clean, Dirty, and Other). Panel (c) shows the year-to-year growth rates of counterfactual shares of Clean and Dirty in all innovations. Panel (d) shows the counterfactual (green) and actual (black) share of Clean relative to the share of Dirty.

Figure 10: Trends re-weighted for Returns distribution



Notes: Trends of Clean and Dirty innovation re-weighted for the private value distribution. Returns are adjusted for grant probability. Panels (a) and (b) show absolute counterfactual trends in Clean and Dirty innovation respectively. The bars (left-hand side y-axis) show the counterfactual counts. The black lines (right-hand side y-axis) correspond to the counterfactual share ($s_{a,t}^{CF(Return)}$) in all innovations (Clean, Dirty, and Other). Panel (c) shows the year-to-year growth rates of counterfactual shares of Clean and Dirty in all innovations. Panel (d) shows the counterfactual (green) and actual (black) share of Clean relative to the share of Dirty.

5 Conclusion

We examine the distribution of private values of clean and dirty innovation. The value of clean innovation tends to be more dispersed and on average more valuable. We also use a new approach to look at the distribution of returns to innovation. In the wake of the Great Recession the distribution of private values as well as returns has become more compressed. This is consistent with the idea that financial restrictions in the wake of the recession have made investors and innovators more risk averse. Because clean innovations tend to be more risky, this could have caused or contributed to the sharp (relative) drop in clean innovation activity after 2011. We quantify the relevance of this by developing a counterfactual methodology that allows us to compute the share of clean (and dirty) innovation that would have prevailed if the distribution of private values (or returns) would

have stayed fixed. This suggests that instead of a clean drop we would have experienced at worst a clean plateau; i.e. the clean share in 2014 would have been no lower than the peak in 2011. We also find that financial constraints have already depressed clean innovation before 2011. As a consequence, in the counterfactual scenario there is still a clean drop but it occurs from a peak that is higher. This suggests that financial constraints have been an important barrier on the path to the clean equilibrium, but that they are not the sole (or indeed the main) explanation for a clean drop in absolute terms.

References

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The environment and directed technical change. *American economic review*, 102(1):131–66.

Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2016). Transition to clean technology. *Journal of political economy*, 124(1):52–104.

Acemoglu, D., Hemous, D., Barrage, L., Aghion, P., et al. (2019). Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution. In *2019 Meeting Papers*, number 1302. Society for Economic Dynamics.

Guillard, C., Martin, R., Mohnen, P., Thomas, C., and Verhoeven, D. (2021). Efficient Industrial Policy for Innovation: Standing on the Shoulders of Hidden Giants. *CEP Discussion Paper Nr. 1813*.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological Innovation, Resource Allocation, and Growth. *The Quarterly Journal of Economics*, 132(2):665–712.

Popp, D., Pless, J., Haščič, I., and Johnstone, N. (2022). Innovation and entrepreneurship in the energy sectory. *NBER Working Paper Nr. 27145*.

Probst, B., Touboul, S., Glachant, M., and Dechezleprêtre, A. (2021). Global Trends in the Innovation and Diffusion of Climate Change Mitigation Technologies. Working paper.

A Details private value estimation

This appendix describes the methodology used to estimate private returns to innovation. To infer the private value of an innovation, we rely on the estimates described in KPSS and an adjustment to account for heterogeneity in grant probability across technological fields. In this section, we give an overview of the event study design KPSS employ to obtain their estimates (we refer to their paper for further details).

Suppose PV_i captures the monopoly rents from exploiting the innovation patented in i. In the absence of any other news, the stock market reaction to the patent grant event is equal to

$$\Delta W_i = (1 - \pi_i) PV_i, \tag{1}$$

where ΔW_i is equal to the difference of a firm's value before and after the moment patent i is granted. π_i is the ex ante probability of any patent being granted conditional on it being public knowledge that the patent application has been made. This expression reflects the assumption that the market knows the value of patent i prior to granting. The day the patent is granted the firm's market value increases by the fraction of the inherent value already known by the market corresponding to the relaxation of the probability that the patent would not be granted. Expression (1) allows to calculate $PV_i = \frac{\Delta W_i}{(1-\pi_i)}$ given an assumption on the ex-ante grant probability. π_i is assumed to be 56% for all patents i, which is the grant rate of US patents between 1991-2001.

This approach to estimating PV_i is subject to the fact that the observed stock market return of any firm might incorporate general movements of the market and unrelated events that might affect stock market returns of the patenting firm. To isolate firm-specific returns that are due to the patent grant, a 'market-adjusted-return model' is used as in **?**. It specifies the firm's idiosyncratic return R_i (i.e. a firm's return around the event minus the return on the market portfolio), as:

$$R_i = v_i + e_i, \tag{2}$$

where v_i is the portion of the return associated to the patent grant event and e_i is the return's component due to unrelated news around the event date. Replacing ΔW_i with the product of the expected value of W_i conditional on the observed R_i and the market capitalization M_i of the firm on the day prior to the event, expression (1) is rewritten as

$$PV_i = (1 - \bar{\pi})^{-1} E[v_i | R_i] M_i. \tag{3}$$

In their preferred specification, KPSS assume a normal distribution for e_i and a normal distribution truncated at zero for v_i. The variance of e_i, as well as the signal-to-noise ratio (the variance of the distribution of v_i divided by the sum of the variances of v_i and e_i) is estimated from the data (the former is allowed to vary by firm; the latter is assumed constant). These parameter estimates allow to calculate private values for a set of $1,801,879$ patent grants published at the USPTO.

Our adjustment takes into account that the probability of a patent grant conditional on application may vary by technological field. 1 in section 2.2 shows that this is the case for the fields we study in this paper. The table is based on all applications filed at the USPTO between 2005 and 2009. This restricted time window allows each application at least 5 years to be granted. We use these grant probabilities to adjust PV_i by plugging in field-specific probabilities for $\bar{\pi}$ for all innovations in our sample. To implement this strategy, we calculate a multiplier to convert original KPSS estimates into adjusted private values as follows:

$$multiplier_i = \frac{(1 - \pi_f)^{-1}}{(1 - \bar{\pi})^{-1}} \tag{4}$$

where π_f corresponds to field-specific grant probabilities for the fields considered in this paper. We then calculated adjusted private values as

$$PV_{i,adj} = multiplier_i \times PV_i \tag{5}$$

B Details cost estimation

In this section, we describe the model developed in Guillard et al. (2021) that is used to estimate the cost of innovating by field. The model imposes structural assumptions on the idea arrival rate and the cost of developing an innovation. This allows to derive theoretical quantile values on the observed distribution of private values of innovations. These quantile values are matched to quantile values observed in the data to obtain estimates of the model parameters.

Assume that a new innovation first requires an idea. Ideas in a given technology class are heterogeneous in quality δ, and follow a Pareto distribution with the following probability density function (pdf):

$$f(\delta) = \begin{cases} \frac{\alpha \mu^{\alpha}}{\delta^{\alpha+1}} & \text{if } \delta > \mu \\ 0 & \text{if otherwise} \end{cases} \tag{6}$$

The support of this quality distribution is $[\mu, \infty)$. α is a parameter that determines the curvature of the idea distribution (with higher values leading to more ideas of low quality). An inventor that has an idea will try to innovate using the idea if it generates private financial gain for her. Her payoff at the time of deciding whether to pursue the idea includes a fixed cost c and also takes into account that the outcome is uncertain. For simplicity, we assume that the probability of innovation success is independent of the idea quality and is a draw from a uniform distribution on the interval $[0, \kappa)$ where $\kappa < 1$. The expected private benefit from innovating conditional on having an idea of quality δ is $PV = \epsilon \times \delta$. Which, because the expected value of ϵ is $\frac{\kappa}{2}$, gives $E\{PV|\delta\} = \frac{\kappa}{2}\delta$.

An inventor chooses to innovate if:

$$E\{PV|\delta\} \geq c$$

Consequently, she will only pursue ideas where $\frac{\kappa}{2}\delta \geq c$. We define λ as the lowest quality idea that will be developed where

$$\lambda = \frac{2c}{\kappa} \tag{7}$$

We are interested in the distribution of idea quality conditional on idea development, which can be written as:

$$f(\delta_i|\delta > \lambda) = \frac{f(\delta_i)}{P(\delta > \lambda)} = \begin{cases} \frac{\alpha\lambda^\alpha}{\delta_i^{\alpha+1}} & \text{if } \delta > \lambda \\ 0 & \text{if otherwise} \end{cases}$$

where $P(\delta > \lambda)$ is the likelihood that any new idea is above the minimum quality required to be developed:

$$P(\delta > \lambda) = \int_\lambda^\infty \frac{\alpha\mu^\alpha}{\delta_i^{\alpha+1}} d\delta_i = \frac{\mu^\alpha}{\lambda^\alpha} \tag{8}$$

We can write the distribution of the private values of ideas that will be developed - i.e. the values we can observe - as:

$$P(PV_i = v|\delta > \lambda) = \int \phi(PV_i = v|\delta) f(\delta|\delta > \lambda) d\delta$$

where

$$\phi(PV_i = v|\delta) = \begin{cases} \frac{1}{\delta\kappa} & \text{if } v < \kappa\delta \\ 0 & \text{if } v > \kappa\delta \end{cases} \tag{9}$$

is the density of PV conditional on δ. Together, these expressions yield:[7]

$$P(PV_i = v|\delta > \lambda) = \int_{\max\{\lambda, \frac{v}{\kappa}\}}^\infty \frac{f(\delta)}{\delta\kappa} d\delta = \int_{\max\{\lambda, \frac{v}{\kappa}\}}^\infty \frac{\alpha\lambda^\alpha}{\kappa\delta^{\alpha+2}} d\delta$$

Consequently

$$P(PV_i = v|\delta > \lambda) = \left[-\frac{\alpha\lambda^\alpha}{(\alpha+1)\kappa\delta^{\alpha+1}} \right]_{\max\{\lambda, \frac{v}{\kappa}\}}^\infty = \begin{cases} \frac{\alpha}{(\alpha+1)\kappa\lambda} & \text{if } \lambda > \frac{v}{\kappa} \\ \frac{\alpha\lambda^\alpha\kappa^\alpha}{(\alpha+1)v^{\alpha+1}} & \text{if } \lambda < \frac{v}{\kappa} \end{cases} = \begin{cases} \frac{\alpha}{(\alpha+1)2c} & \text{if } 2c > v \\ \frac{\alpha 2^\alpha c^\alpha}{(\alpha+1)v^{\alpha+1}} & \text{if } 2c < v \end{cases} \tag{10}$$

where the last equality follows from equation 7.

[7] Note that $f(\delta_i|\delta > \lambda)$=0 if $\delta < \lambda$ from 6. However, we also have that $\phi(PV_i = v|\delta) = 0$ if $\delta < \frac{v}{\kappa}$. This means that $\phi(PV_i = v|\delta) f(\delta|\delta > \lambda)$ will be zero if either of those conditions is binding. This is the reason for the $\max\{\lambda, \frac{v}{\kappa}\}$ expression that is the lower bound of integration.

Notice that the density of PV given $\delta > \lambda$ depends only on c and α. This is because c is a sufficient statistic for the combined effect of κ and μ on the density. Because equation 10 describes the observed innovations , we can estimate parameters α and c by fitting it to the observed distribution of private values PV_i.[8] We can also work out the expected value of the distribution of conditional private values. This is:

$$E\{PV_i|\delta > \lambda\} = \left[\frac{\alpha}{\alpha+1}\frac{v^2}{4c}\right]_0^{2c} + \left[-\frac{\alpha 2^\alpha c^\alpha}{(\alpha+1)(\alpha-1)v^{\alpha-1}}\right]_{2c}^{\infty} = \frac{\alpha c}{\alpha+1} + \frac{\alpha 2c}{(\alpha-1)(\alpha+1)} = \frac{\alpha c}{\alpha-1} \tag{11}$$

The cumulative density is given by:

$$P(PV_i \leq v|\delta > \lambda) = \begin{cases} \frac{\alpha v}{(\alpha+1)2c} & \text{if } 2c > v \\ \frac{\alpha}{(\alpha+1)} + \int_{2c}^{v} \frac{2^\alpha \alpha c^\alpha}{(\alpha+1)w^{\alpha+1}} dw & \text{if } 2c < v \end{cases}$$

We note that

$$\int_{2c}^{v} \frac{2^\alpha \alpha c^\alpha}{(\alpha+1)w^{\alpha+1}} dw = \left[-\frac{2^\alpha c^\alpha}{(\alpha+1)w^\alpha}\right]_{2c}^{v} = \frac{1}{(\alpha+1)} - \frac{2^\alpha c^\alpha}{(\alpha+1)v^\alpha},$$

which means that the cumulative density is

$$P(PV_i \leq v|\delta > \lambda) = \Phi^{PV}(v) = \begin{cases} \frac{\alpha v}{(\alpha+1)2c} & \text{if } 2c > v \\ 1 - \frac{2^\alpha c^\alpha}{(\alpha+1)v^\alpha} & \text{if } 2c < v \end{cases}$$

We can invert this to find quantiles of the distribution. Note that $\Phi^{PV}(2c) = \frac{\alpha}{(\alpha+1)}$. Hence, the p quantile is given by:

$$Q^{PV}(p) = \begin{cases} p\frac{(\alpha+1)2c}{\alpha} & \text{if } \frac{\alpha}{(\alpha+1)} > p \\ \frac{2c}{(\alpha+1)^{\frac{1}{\alpha}}(1-p)^{\frac{1}{\alpha}}} & \text{if } \frac{\alpha}{(\alpha+1)} < p \end{cases}$$

[8] Because 10 won't be differentiable in c and α, we rely on a genetic algorithm to fit the model quantiles to observed ones.

The data give the p-quantile value for every technology class. Hence, we can estimate parameter values for α and c by matching the model quantiles with the data. While this could be done at any level of technology grouping, we implement the estimation at the level of IPC[9] subclasses and year, using estimates for private values developed in KPSS. This results in time-varying estimates for our parameters that can be grouped to any area of innovation by taking a weighted average across innovations belonging to that area.

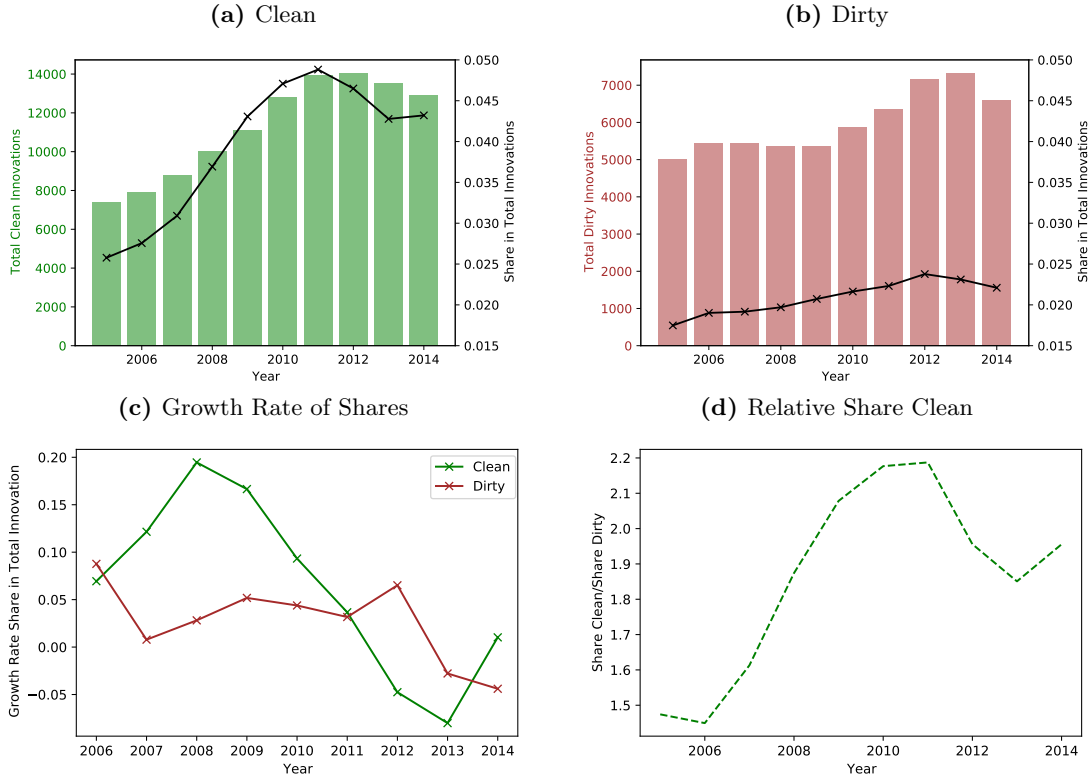
To illustrate our parameter estimation, Figure 1 in section 2.3 shows the modeled and actual distributions for two prevalent IPC subclasses[10] – one with a high and one with a low estimated cost – for the year 2010. The estimated parameter values for this area produce the blue lines in the graphs. The histograms present the actual data.

We use the estimated costs for each IPC subclass to calculate an innovation-specific cost by taking the average across all IPC subclasses a patent is assigned to. To calculate (adjusted) returns, we divide (adjusted) PV_i for innovation i by this weighted average.

[9] For our purposes, the IPC – i.e. International Patent Classification – matches the CPC classification
[10] A61K: 'Preparations for medical, dental, or toilet purposes' and G06F: 'Electric digital data processing'

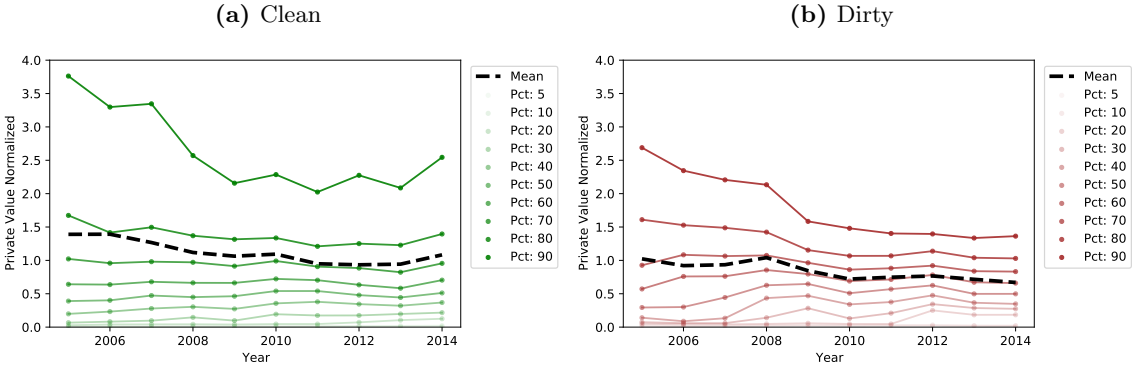
C Additional results

Figure 11: Trends Clean and Dirty – all USPTO



Notes: Trends of Clean and Dirty innovation based on the sample of all inventions for which at least one patent was filed at the USPTO. Panels (a) and (b) show absolute trends in Clean and Dirty innovation respectively. The bars (left-hand side y-axis) show the number of innovations. The black lines (right-hand side y-axis) correspond to the share in all innovations (Clean, Dirty, and Other). Panel (c) shows the year-to-year growth rates of share of Clean and Dirty in all innovations. Panel (d) shows the share of Clean relative to the share of Dirty.

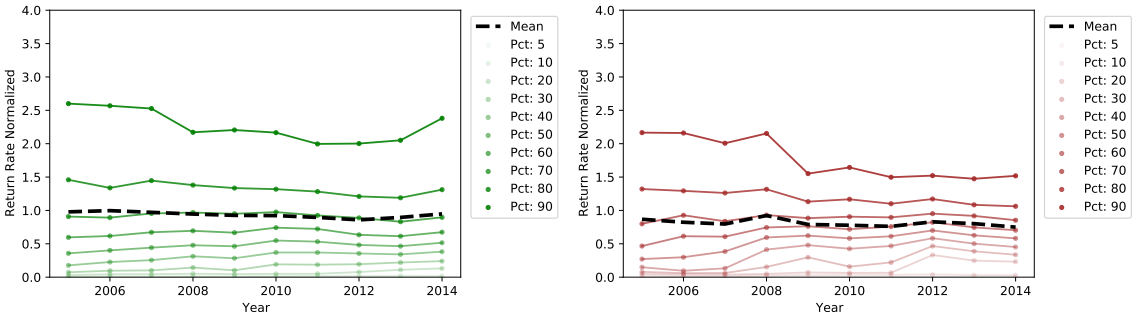
Figure 12: Trends Private Value dispersion



Notes: Evolution of the dispersion of private value in Clean (left) and Dirty (right). Colored lines show the evolution of each decile value (y-axis) between 2005 and 2014. Black line plots the average value for the corresponding category. Private values are normalized by the average private value for the entire sample in each year.

Figure 13: Trends Returns dispersion

(a) Clean **(b)** Dirty



Notes: Evolution of the dispersion of returns in Clean (left) and Dirty (right). Colored lines show the evolution of each decile value (y-axis) between 2005 and 2014. Black line plots the average value for the corresponding category. Returns are normalized by the average return for the entire sample in each year.

Table 2: Counterfactual vs. actual post-2012 innovation

Field	Post-2012 counterfactual increase in innovation (in %)
CCUS	-4.65
Hydrogen	-3.6
Industry	6.44
Biomass & Bioenergy	6.78
Offshore Wind	14.68
Nuclear	22.04
Smart Systems	29.38
Solar	34.92
Tidal Stream	35.84
Building Fabric	36.75
Heating and Cooling	119.69

Notes: