American Economic Journal: Applied Economics 2025, 17(1): 1–40
https://doi.org/10.1257/app.20230052

Do Carbon Offsets Offset Carbon?†

By Raphael Calel, Jonathan Colmer,
Antoine Dechezleprêtre, and Matthieu Glachant*

We develop and implement a new method for identifying wasted subsidies and use it to provide systematic evidence of the misallocation of carbon offsets in the Clean Development Mechanism—the world's largest carbon offset program. Using newly constructed data on the locations and characteristics of over 1,000 wind farms in India, we estimate that at least 52 percent of approved carbon offsets were allocated to projects that would very likely have been built anyway. We estimate that the sale of these offsets to regulated polluters resulted in substantially higher global carbon dioxide emissions. (*JEL* H23, O13, Q42, Q54, Q58)

Carbon offsets have become a popular tool in global efforts to mitigate climate change. These programs work by offering regulated polluters the opportunity to increase their own emissions if they subsidize equivalent emission reductions in unregulated markets. In theory, this allows the same total emissions abatement to be achieved at lower cost. The world's largest carbon offset program—the Clean Development Mechanism (CDM)—has supported more than $90 billion of renewable energy investments in developing countries, equivalent to 13 percent of their total renewable energy investments (Kossoy et al. 2015). By 2030, the CDM will have generated 10.65 billion carbon offsets, roughly equivalent in magnitude to the total emissions of the United States and Europe in 2019.

A key unanswered empirical question is whether these carbon offsets resulted in "additional" emissions reductions in unregulated markets. If carbon offsets are awarded to projects that would have been developed without the subsidy—that is, they are inframarginal—they do not represent emissions savings globally. This results in an inefficient allocation of scarce climate change mitigation resources and higher global emissions. Prior theoretical research highlights the importance

* Calel: Georgetown University (email: raphael.calel@georgetown.edu). Colmer: University of Virginia (email: jonathan.colmer@virginia.edu). Dechezleprêtre: London School of Economics (email: adechezlepretre@gmail.com). Glachant: Mines Paris–PSL. Seema Jayachandran was coeditor for this article. We thank Marika Cabrall and three anonymous referees for their comments and guidance, which have improved the paper. We thank George Akerlof, Antonio Bento, Jishnu Das, Dave Donaldson, Arik Levinson, Molly Lipscomb, David Popp, David Rapson, Lutz Sager, Joseph Shapiro, Jay Shimshack, Sandip Sukhtankar, Alex Teytelboym, John Van Reenen, and Andrew Zeitlin for helpful thoughts, comments, and discussions. We are grateful to seminar participants at American University, the University of Oslo, the University of Virginia, and the Paris School of Economics for their comments and suggestions. Nathan Lado, Eric LaRose, Anna Schröder, and Alex Watkins provided invaluable research assistance. This project was supported by the ESRC Centre for Climate Change Economics and Policy and the Grantham Foundation. All errors and omissions are our own.

† Go to https://doi.org/10.1257/app.20230052 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

of screening out these projects for policy design and welfare (Montero 2000; Fischer 2005; Bushnell 2010; Van Benthem and Kerr 2013; Bento, Kanbur, and Leard 2015; Bento, Ho, and Ramirez-Basora 2015), yet empirical research has struggled to quantify the extent of misallocation due to the difficulty of identifying a credible counterfactual.

This paper makes progress by leveraging the relative simplicity of wind power projects. We propose a new method, which constructs a partial ranking of wind projects based on a few key characteristics, and use it to provide systematic evidence on carbon offset misallocation in the CDM.

Our approach identifies a subset of all inframarginal projects, which we refer to as Blatantly Inframarginal Projects (BLIMPs). BLIMPs are not just inframarginal, but blatantly so, in the sense that they appear strictly more profitable ex ante than other projects that were built without subsidies. The number of BLIMPs provides a lower bound on the degree of subsidy misallocation.

We apply our approach to a new dataset of 1,346 wind power projects in India—a context where it was believed that the CDM could significantly increase development above baseline projections (Purohit and Michaelowa 2007). We independently geo-locate all of the Indian wind power projects identified by Bloomberg New Energy Finance down to the village level (BNEF 2013) and then cross-reference them with individual CDM applications from the UN Environment Program's CDM Pipeline database (UNEP DTU 2021). For each project site, we forecast power output using the hourly distribution of wind speeds and weather conditions from the European Centre for Medium-Range Weather Forecasts' global reanalysis dataset (Muñoz-Sabater et al. 2019). We also estimate the cost of connecting each wind farm to the electrical grid by using detailed grid-infrastructure data from Burlig, Jha, and Preonas (2020). This provides a comprehensive database covering 1,346 Indian wind farms built between 1992 and 2013, of which 476 were registered to receive carbon offsets under the CDM.

Our approach does not require us to estimate the net benefits of each project. This would require strong assumptions and be incredibly data intensive. Instead, we derive sufficient conditions for identifying project dominance based on observable characteristics. We argue that, conditional on being built in the same state and year, a subsidized wind farm strictly dominates an unsubsidized farm if it has a higher capacity, is built in a windier location, and is built closer to a connection point. This is because there are increasing returns to scale in wind farm development because more wind means more power and because both construction costs and transmission losses are increasing functions of distance to a connection point. In this setting, we argue that these three conditions are sufficient to define a BLIMP.

Consider the case of the 91.8-megawatt (MW) CDM-supported wind farm built in Jangi, Gujarat, in 2011. We observe that an unsubsidized wind farm was completed during the same year, just ten miles away in Surajbari. The unsubsidized project had a capacity of only 7.5 MW, was estimated to deliver about 5 percent less power per installed MW because of less favorable wind resources, and was located 3 miles farther away from the nearest electrical substation. These two wind farms were built in the same state and year and were therefore subject to the same policies and market conditions, yet the CDM-supported project was both bigger and better located. We infer that the unsubsidized project at Surajbari was less profitable in expectation, which is a sufficient condition for identifying the larger CDM-supported wind project as a BLIMP.

Out of the 476 CDM-registered Indian wind farms in our data, we identify 257 BLIMPs. For each of these 257 projects, we can point to at least one unsubsidized wind farm built in the same state and year that is strictly less profitable. These projects account for 52 percent of carbon offsets approved for Indian wind projects. Under the assumption that these credits were later used by regulated entities to augment their emissions quotas, global emissions are 28 million tonnes higher, equivalent to running seven coal-fired power plants for a year.[1] The misallocation is so severe that we find that the random assignment of subsidies through a lottery would have resulted in fewer offsets being allocated to BLIMPs.

Our findings are robust to a broad range of sensitivity tests. We perform numerous tests for specific sources of omitted variable bias—by performing comparisons within villages, within turbine manufacturers, within capacity bins, within developer types, and more—and show that a confounding variable would have to be almost perfectly correlated with a project being marginal and with the CDM's observed decisions to undo our findings. Our findings are also robust to a range of possible measurement error models, misspecification tests, and models of spillovers.

Why, then, is the CDM favoring the more profitable projects to such a degree? Possibilities include adverse selection due to large application costs (Chadwick 2006), the manipulation of application documents (Consulate Mumbai 2008; Point Carbon 2010), electoral pressures that obstruct CDM project development (Bayer, Urpelainen, and Xu 2014), routine failures of third-party auditors (Frunza 2013, 63), and conflicts of interest (Transparency International 2011). Our analysis suggests that the problem arises from a series of institutional failures—starting with the Indian authorities explicitly aiming to maximize contributions to its sustainable development goals, which often results in supporting larger wind projects in better locations, and ending with a number of political and resource constraints that limit the CDM Executive Board's ability to independently vet and reject applications forwarded by host countries. This would account for our finding that the CDM has supported even more BLIMPs than would have been supported by a lottery.

Do carbon offsets offset carbon? Our analysis suggests that, in this context, they often don't. Carbon offsets are frequently given to BLIMPs, which means that they do not generate the additional emissions reductions needed to offset the emissions increases that they enable. It is still possible, however, that the remaining CDM-supported wind farms have offset more than their share of emissions. To counterbalance the increase in emissions from subsidizing BLIMPs, we would have to assume that the CDM-supported wind farms that are not BLIMPs reduced India's carbon emissions by 2.12 tonnes per offset. This is equivalent to assuming that the non-BLIMPs are collectively responsible for crowding in half of India's non-CDM

[1] This estimate is taken from the US EPA's Greenhouse Gas Equivalencies Calculator, accessed at https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator (accessed January 13, 2022). The 28 million tonnes of carbon dioxide is also equivalent to putting 6.1 million cars on the road for a year or powering 3.4 million US homes for a year.

wind power capacity.[2] This theoretical possibility doesn't undermine our finding that the CDM has subsidized a large number of inframarginal projects.

If we extrapolate our findings to all CDM projects worldwide and assume no spillovers, we calculate that the program will have approved offsets corresponding to 6.1 billion tonnes of carbon dioxide. These emissions would be equivalent to running 1,500 coal-fired power plants for a year. Applying a social cost of carbon of \$190/tCO_2 (EPA 2022; Rennert et al. 2022), the discounted economic losses associated with these emissions are valued at \$(2020)1.2 trillion.

These findings contribute most directly to our understanding of carbon offset programs. Carbon offsetting is an increasingly important policy tool. A growing number of countries and organizations are now committing themselves to achieving "net zero" carbon emissions with the help of carbon offsets (Black et al. 2021). International negotiations are also underway to develop and implement a successor to the CDM (Michaelowa, Shishlov, and Brescia 2019). The international agreement reached in Glasgow at the COP26 conference in November 2021 promises to grandfather an estimated 3 billion tonnes worth of carbon offsets from the CDM into the new program (Fearnehough, Schneider, and Warnecke 2021). It is critical to consider carefully the design of these programs going forward to ensure that carbon offsets actually offset carbon. Previous work has already identified particular industrial processes that should not be subsidized with carbon offsets (Wara 2007a, b; Schneider 2011; Schneider and Kollmuss 2015), but our results identify a significant misallocation of resources even for the type of projects that the CDM was made to support. Our study suggests that, despite good intentions, the CDM may actually have resulted in higher global emissions. In this regard, our paper connects to a large literature documenting how well-intentioned policies can have unintended consequences (Davis 2008; Oliva 2015; Holland et al. 2016; Agan and Starr 2017; Parker and Vadheim 2017; Bharadwaj, Lakdawala, and Li 2019; Taylor 2019; Doleac and Hansen 2020; Filmer et al. 2023).

Our paper also contributes to the study of a broader class of policies that leverage the logic of offsetting. This includes everything from Corporate Average Fuel Economy standards (Kwoka 1983; Anderson and Sallee 2011; Ito and Sallee 2018) to Renewable Portfolio Standards (Cullen 2013; Gowrisankaran, Reynolds, and Samano 2016; Carley et al. 2018) to key provisions of the Clean Air Act (Shapiro and Walker 2020). Rather than the traditional approach of using tax revenue to subsidize emissions-saving activities, these policies work by encouraging the private sector to cross-subsidize those activities directly without the funds passing through the public treasury. While this approach has some appeal, our findings highlight a potential downside. When the offsets are being generated by an unregulated entity, the consequence of misallocating offsets doesn't merely create an inefficient transfer, as would a traditional subsidy program, but generates external costs by underwriting higher global emissions. When choosing between offset-style regulations

[2] Inframarginal projects may also have spillover effects, crowding in non-CDM capacity. However, any crowding in that arises from inframarginal CDM projects is also inframarginal. These spillover effects would have arisen absent CDM support.

and more traditional subsidy models, the potential external costs from misallocation should be an important consideration.

We also contribute to a broader literature that seeks to identify the existence and magnitude of inframarginal support in subsidy programs. The empirical challenge of identifying inframarginal carbon offsets is similar to the challenges faced by researchers trying to determine whether land owners would have protected ecosystems in the absence of payments (Jayachandran 2013; Jayachandran et al. 2017; Jack and Jayachandran 2019; Jack et al. forthcoming), whether new technologies would have been adopted without rebates (Chandra, Gulati, and Kandlikar 2010; Boomhower and Davis 2014), whether firms would have invested in additional innovation without special tax credits or grants (Hall and Van Reenen 2000; Bloom, Griffith, and Van Reenen 2002; Dechezlepretre et al. 2016; Howell 2017; Azoulay et al. 2019; Pless 2021), or whether firms or workers would have moved in the absence of discretionary incentives (Moretti and Wilson 2014; Slattery 2019; Mast 2020; Slattery and Zidar 2020). In certain cases, researchers have managed to find and exploit discontinuities in rules that have allowed for causal identification. However, this is not the norm. We have explored opportunities to apply these research designs in the context of carbon offsets and are not aware of any successful applications. Our framework offers a different approach to systematically evaluating grant-based subsidy programs.

I. Institutional Background

In this section, we provide relevant institutional details about the CDM, discuss how marginal and inframarginal projects are distinguished in practice through the registration process, and provide background information on the Indian wind power sector.

A. *The Clean Development Mechanism*

The CDM is the largest pollution offset program in the world. It was established as part of the Kyoto Protocol in 1997. Under this program, industrialized countries that had committed to reducing emissions domestically are permitted to meet some of their obligation by developing or financing equivalent emissions reduction projects in developing countries, ostensibly achieving the same global emissions reduction at a lower cost. In practice, the UN issues Certified Emission Reduction (CER) credits to approved projects in developing countries—each signifying one avoided tonne of carbon dioxide—which are then sold to regulated firms in developed countries and counted toward their country's Kyoto target.

From a project developer's perspective, the CDM is much like any other subsidy program. Although developers don't receive money directly from the UN, they do receive valuable CERs that can be exchanged for money on the open market. The promise of this extra money was intended to lure developers to build more renewable energy projects than they would not have built without it.

The CDM has been an extremely popular program. By 2030, it is expected to have generated 10.65 billion offsets, equivalent in magnitude to the total emissions of the United States and Europe in 2019. China (5.94 billion), India (1.23 billion), and Brazil (0.7 billion) account for over 70 percent of these credits. Over half of all CERs finance just two types of projects—hydropower (26 percent) and wind power (23 percent)—the latter being the focus of our study. It has been estimated that the CDM has supported over $90 billion of renewable energy investment in developing countries—roughly 13 percent of their total renewable energy investment (Kossoy et al. 2015). Although the CDM is now no longer accepting new applications, the Paris Agreement promises to expand the program under a new name: the Sustainable Development Mechanism. The follow-on agreement reached at the COP26 conference in November 2021 further commits to grandfathering in 3 billion tonnes worth of carbon offsets from approved CDM projects (Fearnehough, Schneider, and Warnecke 2021).

To register under the CDM, project developers must first submit a Project Design Document (PDD) describing the project and demonstrating "additionality." In this context, additionality means that the project is expected to reduce emissions below the business-as-usual trajectory. The project must then undergo three stages of evaluation. First, a Designated National Authority (DNA)—in India, the Ministry of Environment and Forests—evaluates whether the project, as described in the PDD, meets CDM requirements, which include additionality. Second, a Designated Operational Entity (DOE)—an independent evaluator accredited by the CDM Executive Board—must verify additionality. Third, the CDM Executive Board, which supervises the CDM globally, decides whether to approve registration. If a project is approved, the developer starts receiving CERs as soon as it starts delivering emissions reductions. A hydroelectric dam or wind farm, for instance, would start receiving its approved allotment of CERs as soon as a third party verifies that it has begun generating electricity.

It is critical that CDM projects are marginal, since the CER credits they receive are used to relax another firm's emissions quota. If projects are inframarginal, global emissions increase since no additional reductions have been realized. As in any subsidy program, the regulator's objective is to avoid supporting inframarginal projects.

The CDM Executive Board applies a set of standardized methodologies to determine whether or not a project is marginal. For renewable energy projects, the principle is straightforward: calculate the internal rate of return for the project with and without the extra revenue from selling CER credits. If the internal rate of return with CER revenue exceeds a benchmark rate, but the rate without CER revenue does not, the project is judged to be marginal.[3]

This approach to assessing project additionality is problematic. First, it is difficult to objectively estimate or evaluate internal rates of return. In practice, the PDDs frequently rely on subjective arguments or neglect to provide the underlying data used in their calculations (Schneider 2009). Even when detailed information is available

[3] The CER revenue is estimated by multiplying the electricity that would be generated by a "business-as-usual" emissions factor. Most projects use a factor equal to the generation-weighted average carbon dioxide emissions per unit of net electricity generation from all generating power plants serving the same regional grid (tCO_2/MWh). The assumption is that the new project would replace an equivalent amount of "business-as-usual" generation capacity, avoiding the associated emissions.

and analysis performed, the results are not verifiable by the authorities in charge of CDM registration (Michaelowa and Purohit 2007). Second, even if all the information were correct and verifiable, the methodology itself makes strong assumptions about the growth of renewable power (or more accurately, the lack of growth) in the absence of CER credits.

This raises issues of accountability and regulatory governance. However, exposing the absence of evidence that projects are marginal, as seen in previous studies, is not the same as providing evidence that projects are inframarginal. The most notable direct evidence of inframarginal projects claiming credits concerns the potent greenhouse gas HFC-23, which is a by-product in the production of some refrigerants. Wara (2007a,b), Schneider (2011), and Schneider and Kollmuss (2015) persuasively show that Chinese and Russian refrigerant factories were running overtime to produce more of this by-product since destroying it allowed them to claim CERs that were much more valuable than the refrigerant itself. Once the problematic projects were identified, regulators could easily ban polluters from using these specific credits for compliance purposes.

B. *Wind Power Generation in India*

India accounts for just over 20 percent of all CDM projects and 12 percent of CDM credits, second only to China.[4] The CDM is expected to generate 1.2 billion credits in India by 2030, equivalent to 50 percent of India's carbon dioxide emissions in 2019.

India is estimated to have a total wind power potential between 750 gigawatts (GW) and 1,600 GW, with much of this potential concentrated in the southern and western parts of the country (Phadke, Bharvirkar, and Khangura 2011). At the turn of the century, however, India had almost no wind power generation, and was not building new wind farms. The main barrier to wind farm construction was believed to be the high up-front capital costs (Jagadeesh 2000). Consequently, ex ante evaluations of the Indian wind power sector concluded that there was huge potential for the CDM to finance the construction of additional capacity and maximize the utilization of this untapped wind power potential (Purohit and Michaelowa 2007). This is a setting where the deck is stacked against finding inframarginal projects.[5]

Figure 1 plots the construction of new wind power capacity in India based on Bloomberg's New Energy Finance database (see online Appendix A for details). Within a few years of the CDM's launch in 2000, wind farm construction accelerated and CDM-registered projects began to account for a substantial proportion of new capacity (Figure 1, top panel). The price of CERs had reached about $10 by this point, and it remained high until the final year of the Kyoto Protocol's commitment period. During this period, India experienced a more than twentyfold expansion of

[4] At the outset of this project, we explored the possibility of working in China, but we were not able to identify any wind power projects without CDM support.

[5] Ex ante evaluations do not assess the capacity of regulators to select marginal projects. While Indian regulators may have less capacity than developed countries, it is not obviously more limited than in other major CDM host countries, such as Brazil and China. In addition, while domestic regulators decide which projects to submit for approval, final decisions are made by the centralized CDM Executive Board.

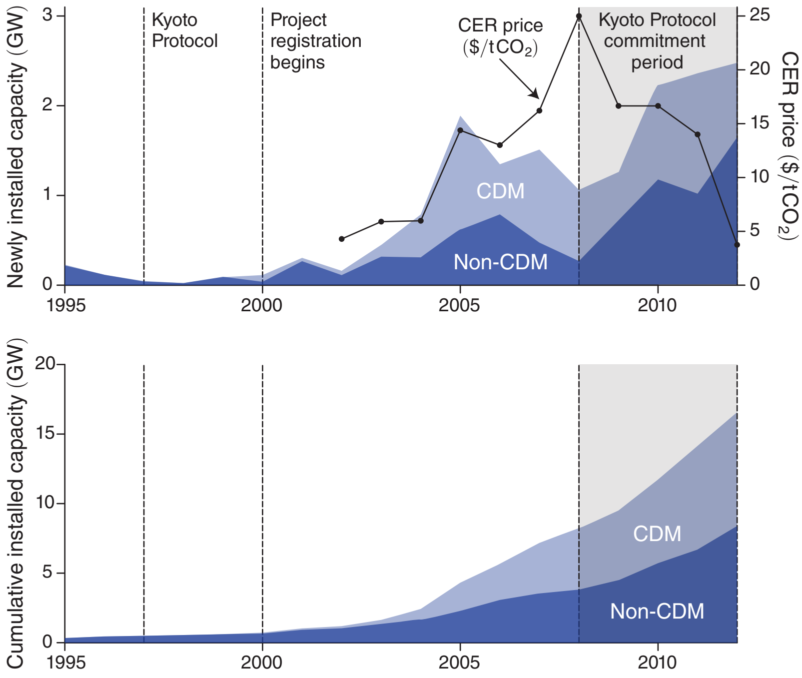

Figure 1. Installed Wind Power Capacity in India

Notes: New capacity is calculated from Bloomberg New Energy Finance and allocated to CDM and non-CDM by cross-referencing projects with the UNEP DTU CDM pipeline. The lower panel plots the cumulative sums of the same data. Annual average CER prices are taken from reports by the World Bank and GIZ (World Bank Group 2019; GIZ 2014). Prices are estimated based on over-the-counter transactions prior to 2008 and then calculated from exchange-based transactions once CERs began to be traded on exchanges.

its installed wind power capacity and today has the fourth-largest installed wind power capacity in the world (Lee et al. 2021). Nearly half of the capacity added between 2000 and 2013 was created through projects registered to receive credits under the CDM (Figure 1, bottom panel). We calculate that these projects were expected to collectively generate roughly 50 million CERs over their lifetime. Our analysis examines the extent to which CDM subsidies contributed to the explosive growth of India's wind power capacity.

The Indian wind power sector was supported by other policies during this period. Most relevant were state-level feed-in tariffs, a national Generation-Based Incentive, and an accelerated depreciation program. As a backdrop for these policies, several states unbundled electricity distribution, creating multiple distribution companies (DISCOMs) that serve different districts within states. These DISCOMs are the main buyers of electricity and differ in their financial condition and credit risk (Ministry of Power, Government of India 2013). We provide more detail about these policies, and how we account for them, in the discussion of our research design below.

By studying the Indian wind power sector, we are consciously selecting a context where the CDM was expected to perform particularly well. This is an important difference from earlier studies, which have tended to look for examples of nonadditionality in places where they are most likely to be found, such as the case of HFC-23 (Wara 2007a,b; Schneider 2011; Schneider and Kollmuss 2015). Such studies are helpful in highlighting a specific type of project that should not be subsidized. By searching for inframarginal projects in a setting where we least expect to find them, we are providing something closer to an upper bound on the CDM's success in identifying and supporting marginal projects.

II. Conceptual Framework

A. *Marginal and Inframarginal Projects*

Following Berry (1992), we consider a simple two-stage game for each market: entry, followed by service provision. In the first stage, each potential entrant, $n = 1, \ldots, N$, makes a choice whether or not to enter market k. In the second stage, entrants make investments that determine postentry payoffs.

Solving backward, we define a payoff function $V(s,x)$ for potential market entrants, where $s \geq 0$ is a subsidy and x is a vector of other payoff determinants. An essential feature of $V(s,x)$ is that it is monotonically increasing in s.

Given the payoff function, potential entrants decide whether or not to enter the market. There is a reservation payoff, R, such that entry occurs if $V \geq R$.

Under these conditions, one can rank potential projects according to their payoffs in the absence of subsidies,

$$(1) \qquad V(s=0,x_1) \geq \ldots \geq V(s=0,x_i) \geq R > V(s=0,x_{i+1}) \geq \ldots \geq V(s=0,x_N),$$

where the subscript indicates the rank order. We assume that V is such that this ranking is preserved for all values of s. Increasing s will only have the effect of shifting some projects from the right side to the left side of R without otherwise disturbing their ranking. Hence, for some subsidy $\bar{s} > 0$, it will be the case that

$$(2) \qquad V(\bar{s},x_1) \geq \ldots \geq V(\bar{s},x_j) \geq R > V(\bar{s},x_{j+1}) \geq \ldots \geq V(\bar{s},x_N),$$

where $j \geq i$. Potential entrants with an index above j (to the right of R in equation (2)) will not enter either with or without the subsidy, so we need not consider them further. Potential entrants with an index between i and j will enter only with a subsidy, which means the subsidy is affecting their decision at the margin. We refer to these as "marginal projects," which in this setting is equivalent to calling them "additional." Potential entrants with an index below i will enter with or without a subsidy. We refer to them as "inframarginal projects," which in this setting is equivalent to calling them "nonadditional."

B. *Blatantly Inframarginal Projects*

The typical approach for determining whether a project, n, is inframarginal is to attempt to estimate $V(0,x_n)$, $V(\bar{s},x_n)$, and R and then check whether $V(\bar{s},x_n) \geq R > V(0,x_n)$. The CDM applies this methodology to evaluate projects ex ante. However, even ex post evaluations using this approach face several challenges, including the fundamentally unobservable nature of $V(0,x_n)$ once a subsidy has been awarded.

Fortunately, one does not need to estimate $V(0,x_n)$, $V(\bar{s},x_n)$, and R to determine whether a project is inframarginal (i.e., whether $n \leq i$). A sufficient condition is that there exists some project $m > n$ that is not receiving a subsidy. Because no projects $m > i$ would enter the market without a subsidy, the existence of an unsubsidized project m is sufficient to infer that n must be less than i.

If one can identify and measure the variables in x and can further identify some subset of these variables that produce monotonic changes in V, which we will denote $\tilde{x} \in x$, then for each project n, one can determine whether or not there exists another project m such that

- m did not receive a subsidy ($s = 0$),
- $\tilde{x}_m \leq \tilde{x}_n$ for each variable in $\tilde{x}$ (where, for convenience, the variables in $\tilde{x}$ are signed so that higher values produce larger values of V), and
- $x_m = x_n$ for each variable not in $\tilde{x}$.

If a project m exists that satisfies the first condition, it must be true that the value of V was large enough to justify entry without the subsidy, which can be stated formally as $V(s = 0,x_m) \geq R$, or equivalently, $m < i$. If project m satisfies the second and third conditions as well, we know that $V(s = 0,x_n) \geq V(s = 0,x_m)$ (or $m < n$). Transitivity implies that $V(s = 0,x_n) \geq R$ (or $n < i$), which means that project n would also have been built in the absence of a subsidy.

If project n is marginal, no other project satisfying all three conditions will exist. If project n is inframarginal, then a project m satisfying all three conditions *may* exist. The existence of a project m is therefore a sufficient condition to infer that n is inframarginal. To distinguish these inframarginal projects from the rest, we refer to them as *blatantly* inframarginal projects, or BLIMPs.

We note that BLIMP designation is an inherently conservative indicator of inframarginality. BLIMPs are likely to be only a subset of inframarginal projects, even when a researcher or regulator is in possession of a complete list of projects. If our list of projects is incomplete, the measure only becomes more conservative. Consider starting with a complete list of projects and then dropping one. If this project is itself a BLIMP, we have reduced the number of BLIMPs by one. If this project is not a BLIMP, its presence could have served to identify other projects as BLIMPs. Any projects that we do not observe will tend to reduce the number of BLIMPs.

This conceptual approach imposes very mild conditions on the payoff function. However, for any given application, it is necessary to specify the payoff function with enough precision that the researcher can determine what variables are contained in x and $\tilde{x}$. We now turn our attention to operationalizing the concept of a BLIMP in the context of wind power generation.

III. BLIMPs in India's Wind Power Sector

In this section, we map the preceding framework onto India's wind power sector to develop an operational definition of a BLIMP in this context. We will discuss the functional form of the payoff function and how we measure the variables that enter into it. The section ends with an empirical test of our identifying assumptions.

A. *An Operational Definition of a BLIMP*

The following equation presents what we consider a reasonable description of the net present value of profits from building a wind farm, n, in state ℓ, in year y.

$$(3)\quad V_{n\ell y} = \sum_{t=y}^{T} \frac{(p_{\ell yt} + s_{nt})(c_n \times f_n)\left[1 - l(d_n)\right] - v_{yt}(c_n) - \tau_{\ell yt}}{(1+r)^{(t-y)}} - F_y(c_n, d_n).$$

The first term is the net present value of the stream of operating profits. Each kilowatt-hour (kWh) of electricity fetches a price of electricity, $p_{\ell yt}$, that may differ across states ℓ, across vintages y, and over time t, as well as a per-kWh subsidy, s_{nt}, that varies across projects and time. Total annual revenue can then be calculated by multiplying the revenue per kWh by the annual power output, which can be written as the product of generation capacity, c_n, and the capacity factor, f_n (the power output per unit of capacity). Total output then needs to be multiplied by a factor that takes account of transmission losses, l. Transmission losses increase as a function of the distance between the wind farm and its closest grid connection point, d_n. The operating profit in year t is what is left over after subtracting maintenance costs, v_{yt}, and taxes, $\tau_{\ell yt}$. Two projects built in the same year, y, with the same capacity, will have the same maintenance cost schedule; if they are built in the same state, they will face the same tax schedule. Finally, the stream of profits is discounted at an annual rate of r, which is common to all projects.

The second term of equation (3) represents the up-front cost of construction, F_y, which depends on the generation capacity, c_n, and distance to a grid connection point, d_n.

As we indicate with our subscripts in equation (3) and will discuss in greater depth shortly, most of the variables vary across but not within states and vintages. Aside from the subsidy rate, only three variables vary across wind farms built in the same state and year, and we argue that V is monotonic in all three: V is an increasing function of generation capacity (c_n) and the capacity factor (f_n) and a decreasing function of the connection distance (d_n). This gives rise to the following operational definition of a BLIMP in the context of India's wind power sector.

DEFINITION 1: *For a CDM-registered wind farm n and an unregistered wind farm m that are built in the same state and year, n is a BLIMP if*

- *it is larger* $(c_n \geq c_m)$, *and*
- *it is built in a windier location* $(f_n \geq f_m)$, *and*
- *it is built closer to a connection point* $(d_n \leq d_m)$.

Three features of this definition are worth noting. First, although equation (3) is helpful as motivation, our empirical analysis does not require us to estimate the value function. We rely only on the milder conditions stated in Definition 1, which are compatible with a larger set of alternative payoff functions.

Second, the definition designates the entire wind farm as a BLIMP when the conditions are met rather than some smaller part of it. This focus on the extensive margin is motivated by the first assumption of increasing returns to scale.[6] For the sake of completeness, we later compute a lower bound on our estimates based on the assumption that BLIMPs are only partially inframarginal.

Third, the identifying assumption is that all other determinants are equal within a given state and year. More precisely, we require that any residual differences are small enough not to affect the ranking of projects. In other words, projects built in the same state and year do not differ significantly in ways that affect their relative payoffs, apart from in their capacity, windiness, and distance to connection points. This assumption is made based on our understanding of the de jure policy and institutional context, but it can be relaxed. In sensitivity analysis we show that our findings are robust to a number of less restrictive "all else equal" assumptions, such as holding the unobserved determinants of profitability constant only across wind farms established in the same village and year.

The remainder of this section motivates the functional form assumptions contained in this definition and describes how we measure each of the variables.

B. *Payoffs Increase in Capacity*

Generation capacity is easy to observe and measure, but its mapping to payoffs is more complicated. Each MW of capacity yields $c_n \times f_n$ kWh of power output in year t, which is then sold at a fixed price of $p_{\ell yt}$. In India, as in many other countries, the price of wind power is agreed in advance through standardized power purchase agreements. Throughout the period of our study, these prices were set by state-wide feed-in tariffs. This means that all wind farms built in the same state and year will earn the same price per unit of electricity (Kathuria, Ray, and Bhangaonkar 2015). These factors suggest constant marginal revenue from each additional MW of capacity.[7]

Marginal revenue could be increasing or decreasing, however, depending on what subsidies are available, s_{nt}. The only relevant domestic program in this period is India's Generation-Based Incentive—a nationwide program that supplements state-level feed-in tariffs with an extra INR 0.5 per kWh of power. The only eligibility criteria is that the wind farm has at least 5 MW built capacity. Given that this supplement is available to all projects above 5 MW, we have no reason to believe that it would serve non-CDM projects disproportionately. We explore the consequences of this assumption in sensitivity analyses.

[6] In other contexts where diminishing returns to scale is more plausible, it would be more appropriate to focus on the intensive margin by designating only a part of the project as blatantly inframarginal.

[7] We caveat that in some states, multiple DISCOMs serve different districts, which could lead to different wind farms in the same state facing different counter-party risk (Ryan 2022). We explore the empirical relevance of this concern below, finding that within-state variation does not meaningfully affect our conclusions.

The only other subsidy to consider is the CDM, which provides registered wind farms with carbon offsets in proportion to their power output. The price of offsets varies over time in response to global demand and supply, but since individual wind farms are price takers in the global offset market, two CDM projects built in the same year will expect to receive the same subsidy per kWh over their lifetime. Non-CDM projects do not receive the subsidy. Unlike the Generation-Based Incentive, eligibility for the CDM is not determined by a capacity threshold. This is why Definition 1 starts by specifying that one wind farm is registered and the other one is not. For such a pair, we can say that the CDM project will have a higher subsidy rate than the non-CDM project.

While the marginal revenue from an additional MW of capacity is constant—or even increasing with size—the marginal cost of building and operating an additional MW of capacity tends to decrease with size. Based on a study of European wind projects during this period, the up-front capital cost of a wind power project, F_y, accounts for as much as 85 percent of the lifetime costs (Blanco 2009). It can be broken down into four categories: turbine cost, construction cost, grid connection cost, and planning cost.

In Europe, the bulk of the up-front investment is the cost of acquiring turbines (65–75 percent) and construction (10–15 percent) (Blanco 2009). We would expect turbines to account for a larger share in India due to relatively lower land and labor costs. We do not observe cost breakdowns for all projects, but for the 69 CDM projects for which we have this information, turbine costs account for an average of 87 percent of the total up-front investment. Both turbine and construction costs scale more or less proportionately with wind farm capacity, which means we should expect a close positive linear association between project costs, F_y, and capacity, c_n. This prediction is supported by our data. From the Bloomberg New Energy Finance project database, we observe both the built capacity and total up-front cost of 30 percent of wind power projects. For this sample, these two variables display a strong positive linear association ($\rho = 0.99$; see also online Appendix Figure D1).

Planning costs are primarily composed of fixed fees that apply to projects of any size, which implies that the function F_y will have a positive intercept. Studies of European wind developments suggest that these costs constitute about 8–10 percent of up-front investment (Blanco 2009). If we regress project costs on capacity and distance for the subset in which all three are observed, the intercept is $\approx$ INR 20 million, which corresponds to nearly 8 percent of project costs on average. The combination of a fixed up-front cost and a constant cost per MW means that the average cost per MW is declining in the scale of the project.

Maintenance costs, v_{yt}, and taxes, $\tau_{\ell yt}$, could in principle undo some of the benefits of scale, but in practice they do not. Maintenance costs are contracted at the time of construction and tend to follow an industry-standard schedule. A typical maintenance contract charges a flat fee per MW of built capacity, which increases by some fixed percentage with each passing year. Thus, two wind farms of the same size and built in the same year will have the same maintenance cost over the lifetime of the project.

When it comes to taxes, wind farms pay the same state and national taxes in India as any other enterprise, with the exception of the Accelerated Depreciation Benefit.

Starting in 1994, wind farm developers were permitted to apply a 100 percent rate of depreciation to their newly built projects. The rate was lowered to 80 percent in 2002, then to 0 percent in 2012, before returning to 80 percent in 2014. This means that wind farm developers in India would have paid almost no taxes on their developments during our period of study. To the extent that wind farm developments are subject to different tax regimes, these will differ primarily by vintage. We have found no evidence of any state tax policies that disproportionately increased the cost of building larger wind farms.

To summarize, power output is an increasing function of generation capacity, as is, to some extent, the revenue per kWh. Because of fixed planning costs, larger wind farms will have a lower average up-front cost per MW. Maintenance and taxes do not increase fast enough—as a function of capacity—to undo the benefits of scale. These facts support the presumption that, at least among wind farms built in a particular state and year, the payoffs to wind farm development are an increasing function of generation capacity. Developers are incentivized to build the largest projects that they are able. In practice, the most important constraints on size appear to be the availability of land and capital.

C. *Payoffs Increase in the Capacity Factor*

For a wind farm of a given capacity, c_n, the amount of power that it generates depends on its capacity factor, f_n. The capacity factor (also known as the plant load factor) is the ratio of actual-to-maximum output. If a 1 MW wind turbine produced at maximum capacity, it would generate 8,766 MWh in a year (1 MW $\times$ 8,766 hours). Since the wind doesn't blow constantly, however, turbines will generate much less power in practice—typically 10 percent to 20 percent of maximum output. The capacity factor measures the windiness of the turbine's location in relation to the optimal wind profile of the turbine.

Under the plausible assumption that windiness itself does not affect building costs, the payoff is an increasing function of the capacity factor. The challenge, in this case, is that neither we nor the developers know in advance how much the wind will blow at a particular site. To assess the value of a potential wind farm, we both have to estimate the capacity factor.

We have already calculated the denominator for a 1 MW capacity factor—simply multiplying the maximum power output by the number of hours in a year—and the same can be done for turbines of any rated capacity, c. To estimate the numerator—the actual power output—we need to know the turbine's power curve:

$$P(\rho, w) = \begin{cases} 0, & \text{if } w < \underline{w} \text{ or } w \geq \bar{w} \\ \min\left\{\frac{1}{2}\rho A w^3 C(w), c\right\}, & \text{if } \underline{w} \leq w < \bar{w} \end{cases}, \tag{4}$$

where the power output P is given as a function of the wind speed, w, and air density, ρ. The swept area, A, the cut-in speed, $\underline{w}$, cut-out speed, $\bar{w}$, power coefficient, C, and rated capacity, c, are all features of the turbine itself.[8]

[8] See online Appendix B for a more detailed description of power curves.

Like many studies working in countries where the quality and quantity of historical weather data are limited, we use reanalysis data to estimate wind resources at each project site (Auffhammer et al. 2013). The ERA5-Land database produced by the European Centre for Medium-Range Weather Forecasts (Muñoz-Sabater et al. 2019) uses a global circulation model to interpolate meteorological variables in observationally sparse regions, yielding data that are more uniform in quality and realism than observations or any model could provide in isolation. ERA5-Land includes a complete set of hourly observations going back to 1981, at a spatial resolution of roughly 9 km over land.[9]

To complete our calculations, our baseline specification uses the technical specifications of the Enercon E-53 800 kW turbine.[10] In our setting, using the same benchmark turbine for all sites has crucial advantages over trying to identify the particular turbines at each location. In particular, when trying to judge which projects are deserving of a subsidy, this approach avoids penalizing developers for choosing superior turbines or rewarding them for selecting inferior ones. Providing a common standard for evaluation avoids these perverse incentives. While there are many different turbines to choose from, the E-53 is by far the most common in our dataset, accounting for about 15 percent of all identified turbines, and it is suitable across a range of locations. It is reasonable to suppose that this turbine was widely available during this period.[11]

D. *Payoffs Decrease in Connection Distance*

Another important consideration for wind farm developers is the cost of connecting the turbines to the electrical grid. According to Blanco (2009), on average, this cost accounts for 10–15 percent of the up-front investment in Europe. The share is likely smaller in India. Based on the 56 CDM project applications for which we can subtract the cost of turbines and land from the total, the remainder is just 11.5 percent on average, which will include connection costs and other miscellaneous costs. Since most of the connection cost is the laying of cables, these costs increase in proportion to the distance between the turbines and the connection point, d_n. Longer connection distances also imply greater transmission losses once the wind farm is built.

We do not observe the connection distance directly, so we must estimate it using the geographical coordinates of the electrical substations and turbines. Coordinates for electrical substations were collected by Burlig, Jha, and Preonas (2020), and coordinates for turbines were found using the Bloomberg New Energy Finance database, the UNFCCC's CDM project database, and extensive online research. Rather than pinpointing turbines individually, we record the coordinates of the village in which the turbines are located. This was partly done out of practical necessity, but

[9] Online Appendix B includes a detailed description of the data and a comparison with historical weather data for India.

[10] See online Appendix B for more details.

[11] Online Appendix B contains a more complete discussion of the choice of benchmark turbine. In robustness checks, we replicate our analysis using an alternative benchmark turbine and using the technical specifications for actual turbines at each wind farm whenever this information is available.

also avoids any possibility that arbitrarily small differences in location could drive our findings. Any variation in location that we might have been able to generate at the subvillage level would likely have contained more noise than signal.[12]

Under the assumption that wind farm developers do their best to minimize connection costs, we calculate the minimum distance to connect each project to the electrical grid using a modified minimum spanning tree algorithm.[13] This approach may underestimate connection distances, since it does not take into account topographical obstacles, but the magnitude of any bias will tend to scale with the connection distance. Since this type of measurement error doesn't distort the *ordering* of connection distances, our empirical analysis is unaffected.

In our setting, a problem would arise only if our estimation method systematically misestimated the distance for either CDM projects or non-CDM projects. This might occur, for example, if CDM projects are disproportionately built where topographical features increase the real connection distance. We find no evidence that there are systematic differences in topographical features between CDM and non-CDM projects (Table 1), and, in sensitivity analysis, we investigate how any differences could theoretically affect our findings by deliberately inflating the distances for CDM projects.

A more remote location could, in principle, lower land costs. However, in 126 CDM applications for which information is available, land costs account for only 2 percent of the up-front investment and, if anything, are weakly positively correlated with our measure of connection distance ($\rho = 0.05$). Given that land costs only account for a small share of overall costs, even substantial differences between CDM and non-CDM projects are unlikely to significantly affect profitability.

The estimated connection distance is positively associated with project costs for the subsample in which both are observed ($\rho = 0.12$; see also online Appendix Figure D2). The association is comparatively weak, consistent with the premise that distance to the grid affects only a small portion of overall project costs. Given the absence of any clear countervailing benefit of remoteness, we assume the payoff is a decreasing function of connection distance.

E. *All Else Equal*

We have argued that, *all else equal*, the developer's payoff is increasing in generation capacity and capacity factor and decreasing in connection distance. We consider the extent to which other factors may vary systematically between CDM and non-CDM projects. We first consider each of the terms that appear in equation (3).

The most important factor is the electricity price, $p_{\ell yt}$, which is set by state-wide feed-in tariffs. Table 1 shows that, if anything, CDM projects tend to benefit from somewhat more generous feed-in tariffs. Table 1 also reports differences in an index

[12] Online Appendix A provides additional details about how the raw data were processed.

[13] To avoid inflating the distances for wind farms spread across widely dispersed sites, we modified the standard algorithm. Instead of trying to connect all points in one step, we first connect the electrical substations to each other. Only then do we extend the graph to include the wind farm sites. This means that the total connection distance will include the edges that connect each cluster of turbines to its nearest substation but avoids any edges that would be necessary to connect distant clusters. See online Appendix C for a more detailed description of the algorithm.

Table 1—Differences in Means for Potential Confounders

	CDM (1)	Non-CDM (2)	Difference (3)	Difference state-year (4)
Electricity prices				
(1) Feed-in tariff (₹/kWh)	3.55	3.52	0.030	–
			(0.6450)	–
(2) Wheeling charge index	0.64	0.60	0.041	–
			(0.2990)	–
(3) DISCOM GPA	3.35	3.40	−0.052	0.002
			(0.4360)	(0.4870)
Credit access				
(4) Bank branch	0.49	0.45	0.037	0.050
			(0.3540)	(0.1220)
Site access				
(5) Paved roads	0.88	0.86	0.014	0.011
			(0.6590)	(0.5600)
(6) Terrain ruggedness index	5.61	5.98	−0.372	0.065
			(0.2580)	(0.6730)
Developer experience				
(7) Prior projects	3.73	3.86	−0.130	−0.861
			(0.8563)	(0.2380)
(8) Prior CDM projects	1.27	0.76	0.515	0.303
			(0.0018)	(0.0710)
(9) Prior non-CDM projects	2.46	3.10	−0.645	−1.164
			(0.2760)	(0.0707)
Timing				
(10) Days after Jan. 1	188.54	251.46	−62.920	−57.846
			(< 0.0001)	(< 0.0001)

Notes: Data on feed-in tariffs and wheeling charges come from Kathuria, Ray, and Bhangaonkar (2015). Data on the DISCOM GPAs come from the Ministry of Power, Government of India (2013). Data on Bank branches, paved roads, and terrain ruggedness come from the SHRUG databases (Asher and Novosad 2023). Data on developer experiences and the timing of projects come from BNEF (2013). Standard errors are clustered at the state-year (rows 1 and 2), DISCOM-year (row 3), village (rows 3–6), developer-year (rows 7–9), and project level (row 10). *p*-values are reported in parentheses.

of wheeling charges from Kathuria, Ray, and Bhangaonkar (2015) (where the most advantageous wheeling charges are coded as ones and the least advantageous as zeros) as well as the Indian Ministry of Power's rating of state DISCOMs operational and financial performance on a scale from A+ to F (translated into a GPA of 4.3 to 0; see Ryan 2022). Higher wheeling charges and lower financial performance of DISCOMs would tend to reduce the effective electricity price below the feed-in tariff level. The differences presented in column 3 are small and statistically insignificant. The difference in DISCOM GPA is 0.002 GPA points once we include state-year fixed effects (column 4).

Maintenance costs, v_{yt}, are relatively small for wind farms. They constitute only 1.6 percent of up-front costs in the 457 CDM applications in which they are reported. Moreover, more than 70 percent of these projects fall within ± 0.5 percent of the corresponding state-year means. As a starting point we assume that maintenance costs are the same for projects built in the same state and year. Tax schedules, $\tau_{\ell yt}$, also do not vary within state-year.

The final piece of equation (3) is the discount rate r. We only observe interest rates for a subset of 227 CDM applications, but the rates quoted do not vary a great deal. The mean interest rate was 11.9 percent, with nearly 75 percent of reported interest rates falling within ± 1 percent of the corresponding state-year means. To get a sense of credit access beyond this set of CDM projects, Table 1 shows that CDM projects are slightly more likely to be built in villages with a bank branch (49 percent versus 45 percent), but the difference is insignificant. This is of course a highly imperfect proxy, but, if anything, it suggests that CDM projects might have somewhat better credit access. CDM projects look to have an even bigger advantage, if anything, when we compare projects built within the same state-year.

The CDM Executive Board evaluates applications by setting V equal to zero and then working out whether the implied internal rate of return $\tilde{r}$ exceeds some common threshold value. This is equivalent to letting r equal some common discount rate and working out whether V exceeds some common reservation payoff. For the purpose of evaluating projects for carbon offsets, then, we argue that a common discount rate should be applied to all projects. Regardless, interest rates do not appear to vary significantly for projects built within the same state and year. Our main analysis therefore assumes that the same discount rate is appropriate for all projects developed in the same state-year. We show that our findings are robust to significantly relaxing this assumption in our sensitivity analysis. [14]

Looking beyond equation (3), we also consider other factors that might cause some projects to be more or less profitable, such as ease of access. One might expect it to be cheaper to build wind farms in more accessible locations, and perhaps CDM projects are disproportionately built in relatively inaccessible locations with high wind speeds, such as more mountainous areas. Table 1 shows two measures of accessibility—whether the village has a paved access road and the terrain ruggedness index from Nunn and Puga (2012). CDM projects are slightly more likely to be sited in villages with paved road access (88 percent versus 86 percent) and in somewhat less rugged terrain (5.61 versus 5.98).[15] The differences are small and statistically insignificant.

Another potentially relevant factor is developer experience. Table 1 shows that CDM projects tend to be built by developers with slightly less experience (3.73 versus 3.86 prior projects). The difference is statistically insignificant, but the breakdown shows a significant difference in the mix of prior projects. New CDM projects tend to be built by developers with more experience building CDM projects in the past, while the non-CDM projects are built by developers with more non-CDM experience. One possible explanation is that the difference captures experience navigating the CDM application process. With no clear way to value the marginal experience, our main analysis assumes that this experience gap does not translate into meaningful differences in project profits. Later, we show that our findings are robust to looking at the subset of developers with at least one CDM project.

Finally, we consider the within-year timing of projects. Costs could vary seasonally, or there may be spillovers from recently completed projects (e.g., positive

[14] For a sufficiently low discount rate, differences in expected project lifetimes, T, could become significant. In practice, there is little variation in expected project lifetimes. The project lifetime is 20 years for 95 percent of the 759 CDM applications where this information is reported. The minimum and maximum are 19 and 25 years, respectively. If the discount rate is anywhere close to the average interest rate of 11.9 percent, these differences in project lifetime are too small to matter.

[15] For scale, the Netherlands has a terrain ruggedness index of 0.037, and Nepal a value of 5.043.

spillovers due to local learning or negative spillovers due to crowding out of available land). Table 1 shows that CDM projects tend to be completed on average two months earlier in the year than non-CDM projects built in the same year. In our main analysis, we assume that within-year differences in timing are not important but later relax this assumption.

Assuming "all else equal" for projects developed in the same year and in the same state is meant as a starting point for our analysis. In subsequent sensitivity analyses, we document the robustness of our findings to less restrictive versions of the "all else equal" assumption.

F. *A Test of the Identifying Assumptions*

Our identifying assumptions are that payoffs are (i) increasing in capacity, (ii) increasing in capacity factor, and (iii) decreasing in connection distance for wind farms built in the same state and year and (iv) that residual differences are small enough not to affect the ranking of projects. In this section we provide evidence in support of these assumptions.

Having linked the set of CDM applications to Indian wind farms, we are left with around 200 successful applications without associated wind farms. One explanation for this is that the wind farm that was eventually built differs so much from the proposed project that we were not able to link it to an application. An alternative explanation is that these proposed projects were never completed because they were unprofitable even with CDM support. In this case, our identifying assumptions imply that the wind farms proposed in these successful applications should not be BLIMPs. If many of these proposed projects are BLIMPs, it would strongly indicate the presence of important unobserved determinants of profitability.

To evaluate this, we first extract information about capacity, location, and timing from these 200 successful applications. Then, treating each proposed project as if it had been realized, we compare it with the set of realized non-CDM projects. We find that only 6.8 percent of these applications would have been considered BLIMPs had they been built. This is substantially lower than the share of BLIMPs among realized CDM projects, as we will show later. This is strong prima facie evidence in support of our identifying assumptions.

IV. How Many BLIMPs Is Too Many?

The ideal number of BLIMPs receiving subsidies is zero. BLIMPs are only a subset of inframarginal projects and thus a conservative indicator of inframarginality. A program could in principle subsidize many inframarginal projects without subsidizing any BLIMPs. The existence of any BLIMPs is a sign that something has gone wrong in the CDM's decision making process.

The number of BLIMPs, however, is a useful indicator of the degree of misallocation. To interpret this number, we need an alternative allocation mechanism of known quality to serve as a benchmark. We can then compare the realized number of BLIMPs in the CDM's allocation to the number of BLIMPs that would have received support in counterfactual allocation scenarios.

We argue that a lottery is a useful benchmark mechanism to assess allocation performance. We randomly select different subsets of the population of wind farms to be awarded CDM credits, noting that, in expectation, a lottery would allocate subsidies in a way that is uncorrelated with whether a project is marginal or inframarginal.

The population of projects might be determined endogenously, however, affecting counterfactual outcomes. First, by holding the population of projects fixed across lottery draws, we are assuming that all non-BLIMP CDM projects are inframarginal. If they were marginal, then they would vanish under some lottery allocations. Removing these projects could cause the number of BLIMPs to rise or fall, on average, depending on the characteristics of non-CDM projects receiving credits in each counterfactual lottery draw. In online Appendix E, we replicate our analysis under the assumption that all non-BLIMP CDM projects are marginal and find that this tends to reduce the number of BLIMPs under counterfactual lottery allocations.

Second, by holding the population of projects fixed across lottery draws, we are assuming that there are no marginal non-CDM projects that the CDM did not fund. If we expand the population of potential projects by adding some number of marginal projects, a lottery will tend to fund those marginal projects at least some of the time. This would reduce the number of BLIMPs in the counterfactual lottery allocation compared to the realized CDM allocation. Therefore, a lottery across existing projects is a conservative benchmark for performance, relative to a lottery over a superset of potential projects.

Third, by holding the population of projects fixed across lottery draws, we are assuming that any marginal CDM projects do not crowd in non-CDM projects. If marginal projects are crowding in non-CDM projects, then a lottery draw that does not include those marginal CDM projects would also have fewer non-CDM projects. If marginal CDM projects are crowding in substantial non-CDM capacity, our analysis could overestimate the number of BLIMPs. In Section VB, we perform several robustness tests but find little evidence that spillovers are important in this context.

By repeating the lottery many times, we obtain a distribution of the number of BLIMPs that a lottery would subsidize. Each time the lottery is repeated, we let the CDM Executive Board randomly select a number of wind farms for registration each year that is equal to the number they actually registered. Some iterations will, by chance, register a lot of small wind farms in remote and windless locations (i.e., few BLIMPs). Other times, the lottery will select a large number of BLIMPs. The distribution gives us a scale. The quality of the realized CDM allocation can be measured by the probability that the lottery results in fewer subsidized BLIMPs. One would hope that no real-world program would produce more subsidized BLIMPs than a lottery.

V. Results

A. *Main Results*

Out of the 1,346 Indian wind farms in our database, 476 are registered under the CDM. Of these, we identify 54 percent as BLIMPs (Figure 2). For each of these 257 projects, we can point to another unsubsidized wind farm built in the same state and year that is smaller, has a lower capacity factor, and is more remote.

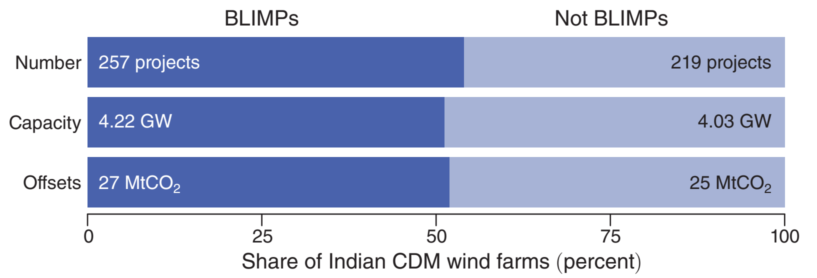

Figure 2. Number, Capacity, and Quantity of CERs Approved for BLIMPs

The CDM appears to have performed poorly relative to the allocations that a lottery mechanism might have produced. Figure 3, panel A plots the observed number of BLIMPs against the distribution of 100,000 hypothetical lottery realizations. The *p*-value tells us that, for two out of every three random draws, a lottery would give rise to fewer BLIMPs than the CDM. Had the CDM chosen projects at random from the whole population of wind farms, it would likely have registered fewer BLIMPs.

Not all BLIMPs represent mistakes of equal magnitude. It is possible that the CDM's registration process gives rise to a relatively large number of small BLIMPs, whereas the imagined lottery might be consistently distributing carbon credits to a smaller number of much larger wind projects. We explore the severity of misallocation by looking at the total capacity of BLIMP projects.

Figure 2 shows that the BLIMPs registered under the CDM have a collective capacity of 4.22 GW, accounting for about a quarter of India's total installed wind capacity in 2013. Figure 3, panel B shows that 980 times out of a 1,000, a lottery would have subsidized less BLIMP capacity than the CDM.

Ultimately, we care about the number of BLIMP carbon offsets, not the number or capacity of the BLIMPs themselves. This requires an extra step. Unlike the number or capacity of projects, we do not observe the number of CERs that would have been given to every wind farm in our data. We need this number to determine how many CERs the lottery would have awarded under counterfactual scenarios. Fortunately, the quantity of CERs received by CDM projects is strongly predicted by capacity (Figure 4). We use this relationship to impute the quantity of CERs that a non-CDM project would be awarded under alternative allocation scenarios.

Figure 2 shows that out of the 53 million CERs estimated to have been issued to Indian wind farms, just over half—about 27 million—went to BLIMPs.[16] When sold as offsets, these carbon credits have allowed regulated polluters around the world to emit more than 27 million additional tonnes of carbon dioxide. Since a tonne of coal produces around 2.25 tonnes of carbon dioxide when burned, these carbon credits

[16] The CDM database lacks information on the number of offsets issued for about half of our wind farms, generally smaller ones. We have used the linear fit in Figure 4 to impute their offset quantities—or at least what would have been reasonable for project developers to expect. For comparison, the maximally conservative approach would be to impute zero offsets wherever the values are missing. This returns an estimate of 41 million issued CERs, of which 21 million went to BLIMPs.

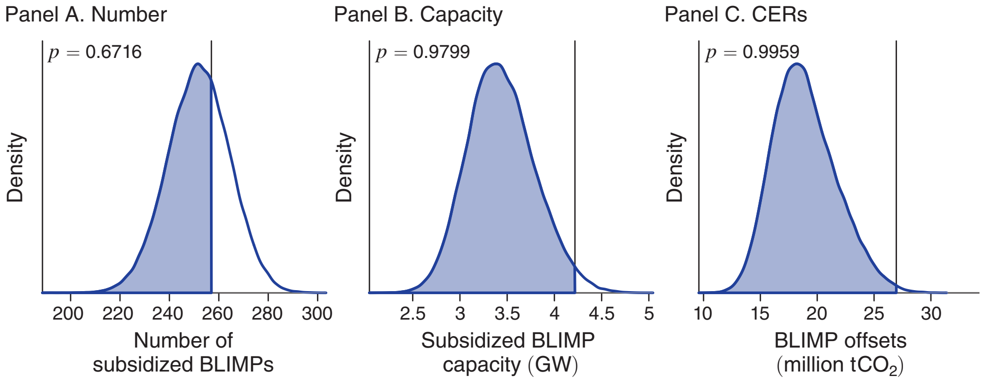

Figure 3. Comparing the Number, Capacity, and Quantity of CERs to a Lottery

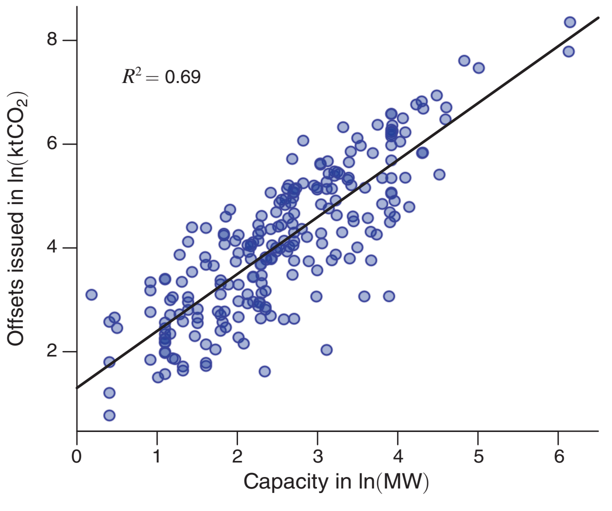

Figure 4. The Relationship between Capacity and Issued Carbon Credits

have given license to burn an additional 12 million tonnes of coal—enough to run seven coal-fired power plants for a year.[17] Figure 3, panel C shows that a lottery would have awarded fewer carbon credits to BLIMPs 996 times out of 1,000.

Our estimates are derived under the assumption that, aside from capacity, windiness, and remoteness, all else is equal for two projects built in the same state and year (see Section IIIE). If CDM and non-CDM projects built in each state and year vary systematically in unobserved ways, this may undermine our analysis. For any omitted variable to be problematic, it must have two features: (i) it increases a project's chances of becoming registered under the CDM, *and* (ii) it increases a project's chances of being classified as a BLIMP. For example, CDM officials might observe

[17] If we extrapolate this rate of inframarginal support to the CDM as a whole, the program would have allowed regulated polluters to emit an additional 6.1 billion tonnes of carbon dioxide. This is equivalent to running approximately 1,500 coal-fired power plants for a year.

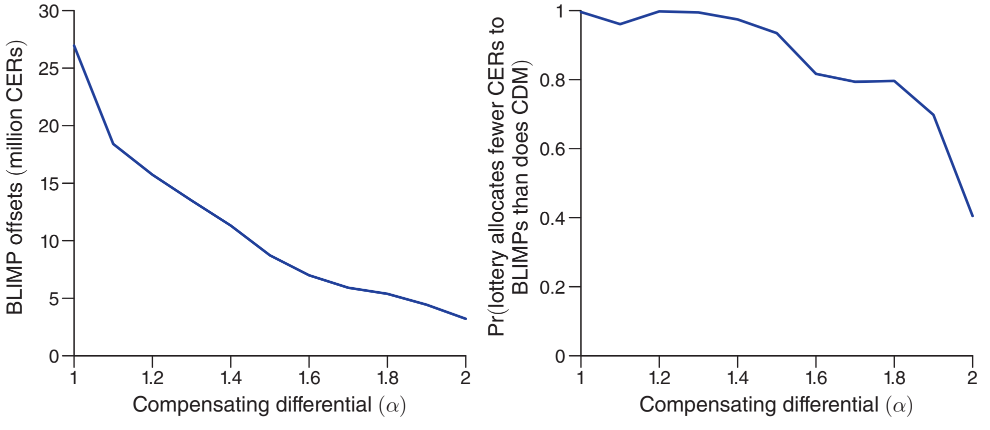

Figure 5. Sensitivity to Unobserved Costs, Measured in Terms of a Compensating Differential in Capacity, Capacity Factor, and Connection Distance

some project cost (hidden from us) that supports the claim of it being marginal. If those projects are bigger or in better locations to partially compensate for the hidden cost, our analysis might inappropriately classify them as BLIMPs.

To evaluate the consequence of this kind of omitted variable we adopt a stricter definition of BLIMPs. We define a subsidized project as a BLIMP only if another project built in the same state and year has at least $\alpha \geq 1$ times the connection distance, at most $1/\alpha$ times the capacity, and at most $1/\alpha$ times the capacity factor. The scalar α measures the hidden cost as a compensating differential on observables. In the main analysis, $\alpha = 1$. Figure 5 shows what happens when we increase the value of α. For $\alpha = 1.2$, the number of BLIMP offsets falls to 16 million. For $\alpha = 2$, we are demanding to see non-CDM projects that are twice as remote, with half the capacity, and with half the wind before calling something a BLIMP. The number of BLIMP offsets declines even further, as expected, but remains above 5 million CERs. The CDM's performance does not, however, improve much compared to a lottery. This is because the stricter definition also reduces the number of BLIMPs under counterfactual assignments.

A second way to understand the consequences of an omitted variable is to change the counterfactual. Our main lottery analysis assumes that, conditional on state and year, the probability of becoming a CDM project is the same whether the project has a high or low value of the omitted variable. Denote this case as $\Gamma = 1$. We could instead posit that a high value of the omitted variable doubles the chances of registering under the CDM ($\Gamma = 2$), or triples it ($\Gamma = 3$), quadruples it ($\Gamma = 4$), and so on.[18] This amounts to saying that CDM officials have access to a more and more informative signal of project marginality—a signal that we cannot observe. Ratcheting up Γ won't affect the number of BLIMPs, since the definition of a

[18] This type of sensitivity analysis is mathematically identical to Rosenbaum and Silber (2009).

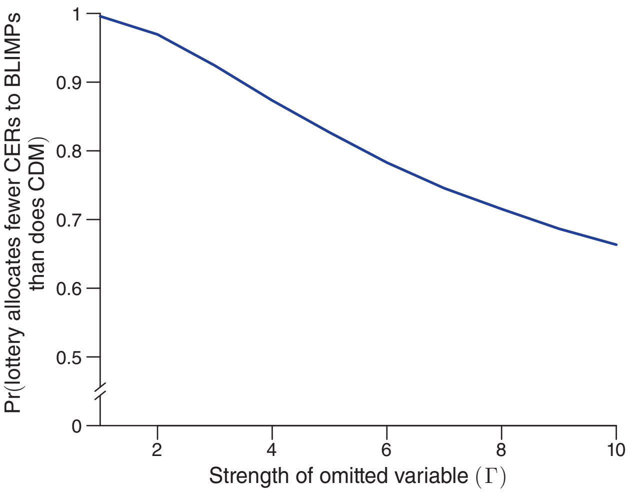

Figure 6. Sensitivity to a Generic Omitted Variable

BLIMP isn't changing, but the counterfactual allocations will shift closer and closer to the realized CDM allocation, which will reduce the p-values.

Figure 6 shows what happens to the p-value as we increase Γ all the way up to 10, at which point we are positing an omitted variable that produces a tenfold increase in the odds of registering under the CDM. Increasing Γ in this way does reduce the p-value substantially, but even at the top end, a lottery still assigns fewer CERs to BLIMPs 70 percent of the time.

In summary, despite the Indian wind power sector having been identified ex ante as offering significant opportunities to support marginal projects, we find evidence of meaningful support for inframarginal projects. Our analysis reveals that the program registered a large number of BLIMPs, performing poorly compared to a lottery. We find that wind farms supported by the CDM tend to be larger and receive more carbon credits, on average, than the BLIMPs that would have been subsidized under a lottery. This central conclusion is robust even in the face of very substantial hypothetical omitted variable biases. Our results pose a serious challenge to claims that the CDM is successfully offsetting emissions. Collectively, the BLIMPs were approved to receive 27 million CERs, despite being very unlikely to have reduced emissions.

B. *Sensitivity Analysis*

This section investigates a number of possible alternative explanations for our findings. We begin with more quotidian explanations such as measurement error and misspecification and then proceed to more exotic explanations like spillovers and partial inframarginality.

Measurement Error.—Our analysis relies on two inputs that have been estimated—connection distances and capacity factors. These may have been measured with error.

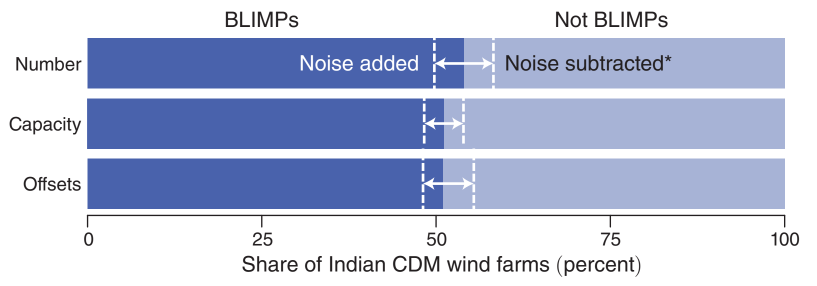

Figure 7. Classical Measurement Error

Notes: The dashed white lines to the left show the number of BLIMPs, their capacity, and their offsets after adding mean zero normally distributed noise to our estimates of connection distances and capacity factors (mean of 1,000 replications). If the effect of classical measurement error is locally well-behaved, we can infer that the number of BLIMPs, their capacity, and their offsets would be correspondingly greater if we were able to subtract any hypothetical measurement error from our data. This is illustrated by the dashed white lines to the right. The asterisk indicates that these lines are inferred rather than the result of direct simulation.

Classical measurement error can result in upward or downward bias in this setting. The bias depends on the true fraction of CDM projects that are BLIMPs and on the relative likelihood of misclassifying BLIMPs and non-BLIMPs in the presence of measurement error. Consider the case in which mismeasurement of project characteristics leads us to misclassify BLIMPs and non-BLIMPs with equal probability. If more than half of the CDM projects are BLIMPs, we would misclassify a larger absolute number of BLIMPs as non-BLIMPs. Measurement error leads to a downward bias in this scenario. If fewer than half of the CDM projects are BLIMPs, we would expect an upward bias. The misclassification probabilities, however, need not be equal for BLIMPs and non-BLIMPs. These probabilities depend on the distribution of project characteristics. Even when more than half of the CDM projects are BLIMPs, there might be an upward bias if BLIMPs are less likely to be misclassified than non-BLIMPs. By the same logic, if BLIMPs are more likely to be misclassified than non-BLIMPs, there could be a downward bias even when fewer than half of the CDM projects are BLIMPs.

We cannot sign this hypothetical bias analytically. We can sign it empirically, however, by adding a small amount of noise to our data and seeing how it affects the number of BLIMPs. Figure 7 shows that the estimated number of BLIMPs tends to shrink when we add a small amount of normally distributed noise. This suggests that, if classical measurement error ware already baked into our data, our original estimates may understate the true number of BLIMPs.

To investigate whether we have introduced bias through nonclassical measurement error, we start by recomputing our results using alternative estimates of connection distances and capacity factors. Since each method of estimation is likely to produce different errors, it would be revealing if this exercise yielded substantially different results.

One potential source of nonclassical error in our estimates of connection costs is our assumption that wind farms are always connected to the nearest substation. It might be that CDM projects need extra support because they face greater obstacles connecting to nearby substations and must therefore connect to more distant substations. In this case, we will have systematically underestimated the connection costs of CDM projects relative to non-CDM projects. To address this concern, we re-estimate connection distances while imposing the constraint that wind farms may only be connected to substations in the same state. This should redress any systematic imbalance between CDM and non-CDM projects that might arise from differing administrative obstacles associated with transecting state boundaries. The results, reported in row 1 of Table 2, are almost identical to before. We report all of the different outcomes for completeness, but our main interest is the number of offsets allocated to BLIMPs, which, if anything, is slightly larger. The CDM's performance relative to a lottery, measured by the bracketed value below, is also practically unchanged.

Another possible source of error in our estimates of connection distances is the comparative difficulty of obtaining high-quality data on the locations of electrical substations. Our main analysis uses a dataset compiled by Burlig, Jha, and Preonas (2020). If this list happens to miss more substations without nearby CDM projects, then we would artificially increase connection distances for non-CDM projects. We can examine this hypothesis indirectly by substituting a range of alternative lists of plausible grid connection points. In row 2 of Table 2, we report the results using the locations of conventional power plants. In row 3, we use the location of cities that, according to the 2001 Indian census, had a population of at least 100,000. In row 4, we use the location of cities listed in the 2001 Indian census as having electrical power. The results vary between 23 and 30 million BLIMP offsets, compared to our main estimate of 27 million, indicating that either all four lists suffer from the same bias or the list compiled by Burlig, Jha, and Preonas (2020) does not suffer from systematic omissions.

Turning to our estimates of capacity factors, we consider two possible sources of measurement error—air density estimation and the benchmark turbine. We have estimated capacity factors using available weather data on air density, but if our estimates of air density are noisier for some locations than others and this pattern correlates with CDM registration, it could affect our results. Row 5 of Table 2 shows that using a common air density has no meaningful effect on our findings.

The choice of benchmark turbine could similarly introduce nonclassical measurement error. In our main analysis we use the technical specifications of Enercon's E-53 turbine to estimate capacity factors. This is the most common turbine in our database. The power curve of any particular turbine, however, will undoubtedly favor some wind profiles over others. If these locations happen to be correlated with CDM registration, this would bias our results. We address this concern by swapping out Enercon's E-53 turbine for Suzlon's S82 turbine—another very common turbine in our dataset that has a different power curve. Row 6 of Table 2 shows that this substitution makes very little difference in the results.

To address the issue more generally, we ask how much nonclassical measurement error would be necessary to qualitatively alter our findings. If connection costs are systematically underestimated for non-CDM projects compared to CDM projects—or capacity factors systematically overestimated—we would need to apply an adjustment factor before comparing projects. To explore this, we impose the assumption that CDM projects are actually $\beta \geq 1$ times more remote and only

Table 2—Summary of Sensitivity Analyses (p-Values in Parentheses)

	BLIMP share (in percent)	BLIMP capacity (in GW)	BLIMP offsets (in million tCO_2)
Main result	54	4.218	26.957
	(0.6716)	(0.9799)	(0.9959)
Measurement errors			
(1) Connect within states	55	4.314	27.277
	(0.7463)	(0.9875)	(0.9967)
(2) Connect to power stations	51	4.803	30.418
	(0.7062)	(0.9997)	(0.9999)
(3) Connect to cities of $>$100,000	54	4.145	24.905
	(0.9143)	(0.9993)	(0.9999)
(4) Connect to cities with power	50	3.899	23.577
	(0.7176)	(0.9874)	(0.9936)
(5) Standard air density	54	4.218	26.957
	(0.6693)	(0.9798)	(0.9959)
(6) Suzlon benchmark turbine	54	4.194	26.774
	(0.6090)	(0.9781)	(0.9953)
Misspecification tests			
(7) Match connection distance	23	1.661	10.151
	(0.0449)	(0.9509)	(0.9747)
(8) Match capacity factor	20	1.444	8.229
	(0.0386)	(0.9573)	(0.9853)
(9) Match capacity	10	0.486	2.513
	(0.0012)	(0.9688)	(0.9801)
Omitted variables			
(10) Match manufacturer	27	1.382	6.476
	(0.0850)	(0.5882)	(0.5878)
(11) Turbine selections bound	46	2.662	15.445
	(0.2513)	(0.8416)	(0.9288)
(12) Match number of sites	37	2.752	15.022
	(0.0407)	(0.8309)	(0.8393)
(13) With 5 MW threshold	42	3.161	18.087
	(0.2388)	(0.7931)	(0.8309)
(14) Within district-year	30	2.399	12.789
	(0.0351)	(0.8421)	(0.8234)
(15) Within village-year	13	0.809	4.229
	(0.0007)	(0.6773)	(0.7013)
(16) CDM developers only	34	2.870	17.476
	(0.3284)	(0.5544)	(0.7005)
(17) Timing	54	4.244	26.703
	(0.7908)	(0.9871)	(0.9951)
Spillovers			
(18) Without financial spillovers within developers	44	3.532	22.481
	(0.3813)	(0.8770)	(0.9534)
Allowing for mistakes			
(19) Two inferior projects	31	2.524	14.426
	(0.0613)	(0.7107)	(0.7957)
Partial inframarginality			
(20) Next biggest project bound	54	1.899	8.321
	(0.6716)	(0.6094)	(0.2800)

$1/\beta$ times as windy as our estimates indicate. As we increase β, CDM projects look less and less desirable as investment opportunities compared to non-CDM projects.

A β of 1.2, for instance, inflates connection distances of CDM projects by 20 percent and reduces their capacity factors by 20 percent. With this more stringent definition of a BLIMP, the number of BLIMP offsets drops to about 17 million, but the CDM still just barely outperforms the lottery mechanism. If we impose an even larger handicap on CDM projects, the CDM begins to outperform the lottery. Online Appendix Figure D3 shows how the number of BLIMP offsets and the probability of outperforming the lottery fall as β rises from 1 to 2.

Misspecification Tests.—Our empirical analysis relies on the argument that, during this time period in India, larger wind farms with higher capacity factors that are located closer to electrical substations are expected to be more profitable developments. It is worth investigating which of these monotonicity conditions is most important in determining the result.

In rows 7, 8, and 9 of Table 2, we relax each monotonicity condition in turn. We accomplish this by making sure that CDM wind farms are first matched to non-CDM wind farms based on either connection distance, capacity factor, or capacity and then use the remaining two monotonicity conditions to determine whether the project is a BLIMP. In row 7, we match CDM wind farms to non-CDM wind farms within ± 5 percent of their connection distance, d_n, and then use their relative capacities and capacity factors to determine whether each project is a BLIMP. For rows 8 and 9, we match on capacity factor or capacity and rank order projects based on the remaining two monotonicity conditions.[19]

Matching projects along more dimensions naturally results in a drop in the size of the test statistics. Still, the CDM is more likely to produce greater BLIMP capacity and a greater quantity of BLIMP offsets than a lottery. We also see that most of the identification of BLIMPs in our original analysis comes from relative capacity. This is reassuring for two reasons. First, of the three variables that enter monotonically into the value function, capacity is the only one for which we have direct measurements, reducing the risk that our results are due to measurement error. The second reason is the strong monotonic association between capacity and project costs in our data (see online Appendix Figure D1).

Omitted Variables.—In Section VA, we showed that our conclusions are relatively insensitive to omitted variable biases. In this section, we investigate more directly how specific omitted variables, many of which were discussed in Section IIIE, might create bias.

One possible omitted variable is the turbine a project developer selects. To the extent that different manufacturers differ in their pricing, quality, technical specifications, ability to supply large developments, maintenance costs, etc., the choice of turbine could affect profitability in ways not fully captured by a wind farm's capacity and capacity factor. This would be missed in our original analysis, which compares across CDM and non-CDM projects with different turbine suppliers. To explore the contribution of turbine selection we use the information we have on a subset of projects to restrict comparisons to projects that use turbines from the same manufacturer (see row 10 of Table 2). Unsurprisingly, BLIMPs are much rarer if we

[19] For capacity factor and capacity, we were able to match projects within ± 1 percent without totally negating our ability to identify any BLIMPs.

require that wind farms be built in the same year and same state and use turbines from the same manufacturer. However, taking into account that the same difficulty extends to counterfactual CDM assignments, we see that the CDM is still expected to allocate more offsets to BLIMPs than the lottery mechanism.

Another way to evaluate the role of turbine selections is to reestimate capacity factors using each project's actual turbine, rather than a common benchmark.[20] For this we have to rely on a subset of projects for which we have been able to identify the turbine make and model and then track down the relevant technical specifications.[21] In row 11 of Table 2, we see that the number of BLIMP offsets is about 15 million when considering only this subset. A lottery is expected to outperform this result more than 90 percent of the time. If we were to impute the very worst technology to the remaining CDM projects and the best technology to the remaining non-CDM projects, this would result in, at best, zero additional BLIMPs—the number of BLIMP offsets could go up if we were to have more information, but it could not fall. The 15 million BLIMP offsets is therefore a lower bound with omitted information on turbine selections.

Another potential omitted variable is the number of sites a project spans. Like the turbine selection, the number of sites could affect project costs in ways not captured by capacity or connection distance. Row 12 matches projects on state, year, and number of sites. As before, this makes it harder to identify BLIMPs, but the CDM still performs poorly compared to a lottery.

Another possibility is that projects above 5 MW are somehow different from projects below 5 MW. As mentioned earlier, the Generation-Based Incentive specifically supports wind farm developments that exceed 5 MW in capacity. If the policy is in place to compensate for some unobserved cost of scale, such as the loss of support from local governments, then our results may be biased. In row 13, we match wind farms on the year of construction, the state, and whether they exceeded the 5 MW capacity threshold. This improves outcomes for the CDM, but a lottery would still perform better most of the time.

Another class of omitted variables are those that vary spatially within states. We considered some of these in Table 1, but the list is theoretically infinite. Perhaps there is systematic variation in some unobserved policies or factor prices between districts or villages within the same state that affect costs (e.g., different DISCOMs). Alternatively, the CDM may systematically favor developments in certain areas for reasons other than emissions reduction potential (e.g., less developed regions, regions with worse air quality, or regions where the electricity grid is being expanded).

[20] To be precise, we have to replace capacity factors with power per acre for this exercise to enable comparisons across turbine technologies. The capacity factor measures the fraction of potential output, but turbines differ in potential output. In particular, larger turbines typically need to be spaced further apart, and they are typically mounted on taller towers that expose them to greater wind speeds. To account for these differences, we estimate the power output that each turbine configuration would yield per acre in each wind farm location. See online Appendix B.1 for details on these adjustments. We continue to refer to capacity factors in the text for convenience.

[21] We purchased a proprietary dataset with technical details of 980 turbines catalogued as of 2014. After excluding offshore turbines, turbines with missing technical information, turbines from manufacturers that never appear in our data, and turbines listed as "under development" in 2014, we were left with 54 different turbines, available in over 100 configurations (e.g., turbines mounted on towers of different heights). After complementing this with additional manual searches, we could link turbine specifications to 673 projects in our database.

To address this, rows 14 and 15 show the results of matching wind farms built in the same district-year and the same village-year, respectively. This eliminates the influence of any unobserved policies or factor prices that vary systematically across districts or villages within a state.

The number of BLIMP offsets falls substantially in these two specifications, though this is mainly due to the fact that many districts and villages do not have both CDM and non-CDM wind farms. Matching at the village-year level has the side effect of also eliminating variation in connection distances and capacity factors, so BLIMPs are only identified based on size differences. Even so, measured either by BLIMP capacity or offsets, the CDM still performs poorly relative to a lottery.

It is also possible that differences in developer characteristics matter. Project developers may differ systematically in their engagement with local stakeholders, in their adherence to local regulations, and in their ability to develop wind projects more broadly. For example, the CDM may only support projects put forward by reputable developers. If smaller, more remote wind farms are built by less reputable developers, this would lead to the CDM supporting larger projects. If reputation is observed by the CDM, this could drive our results. Row 16 shows that the CDM's performance relative to a lottery improves when we limit our sample to projects built by developers with at least one previous CDM-registered wind farm. Even within this restricted sample, the lottery outperforms the CDM roughly 700 times out of 1,000. While there might be some merit in the "developer quality" hypothesis, it explains at best a small part of the CDM's apparently poor performance.

Finally, Table 1 shows that CDM projects tend to be built earlier in the year than non-CDM projects. It is conceivable that these projects are marginal at the time of construction, but, owing to favorable economic and technological trends, they often achieve private profitability within a few months. Projects may appear to be BLIMPs to us if the CDM is accelerating development by a few months. To address this, row 17 of Table 2 compares projects only to others built in the preceding 12 months, rather than matching on calendar year. The results are practically identical.

In summary, we have considered several important classes of omitted variables—ones that vary with turbine selections, with project fragmentation, with capacity, with timing, with developer characteristics, and across space. Our conclusions from these robustness tests generalize to any omitted variables that correlate strongly with the variables we have considered. Naturally, the number of BLIMPs declines whenever we apply more restrictive criteria, but ultimately none of these omitted variables are sufficient to eliminate the problem of BLIMPs. In each case, the CDM continues to perform poorly compared to the lottery benchmark.

Spillovers.—Row 17, which restricts comparisons to non-CDM projects completed in the prior 12 months, also tests for the presence of local spillovers. If building a CDM project creates more favorable local conditions for development of other wind farms—local learning, permitting, etc.—then our main estimate might misclassify it as a BLIMP because of subsequent projects it enabled. By comparing only to others built in the preceding 12 months, row 17 shows that this spillover channel does not account for our findings.

More broadly, since inframarginal projects would have been built without a subsidy, any spillover effects that are generated would also have occurred without the subsidy. Inframarginal projects generate inframarginal spillovers. The exception to this would be if the subsidy transfer itself creates a spillover. In that case, even a subsidy to an inframarginal project may have some benefit. These financial spillovers are most plausible for projects built by the same developer. A developer may choose to put their best projects forward to the CDM, instead of their worst projects, but may nevertheless use the offsets to cross-subsidize marginal projects. In the presence of such financial spillovers, it would not always be possible to infer that a project built without direct support from the CDM is in fact inframarginal.

To investigate the empirical relevance of this financial spillover mechanism, we remove from consideration all non-CDM projects that were built after their developers received support through the CDM. Although these projects appear unsupported to us, it is possible that they are marginal projects that were indirectly subsidized by their developers using offset revenue from the CDM. Removing them from consideration means that we will not identify a BLIMP using a non-CDM project that could have received support indirectly. Row 18 of Table 2 shows that we still recognize almost half of the CDM projects as BLIMPs, collectively accounting for over 22 million offsets. Randomly assigning CDM support through a lottery would be expected to award fewer offsets to BLIMPs more than 95 percent of the time. Within-developer financial spillovers therefore do not appear to be an important explanatory factor in this context.

At an even broader level, it is possible that, by supporting some marginal projects, the CDM may have indirectly helped to establish the wind power sector in India, providing market-level spillovers. These spillovers would not affect our estimates of BLIMPs, since these benefits would accrue to future CDM and non-CDM projects alike. However, market-level spillovers from marginal projects could alter the aggregate effect of the CDM on emissions. If marginal projects crowd in additional renewable energy capacity, each carbon credit awarded to them would potentially offset more than one tonne of carbon dioxide.

Assuming that all non-BLIMP CDM projects are marginal, we calculate that each offset awarded to a non-BLIMP would have had to generate an average emissions reduction of 2.12 tonnes to offset the increase in emissions from the CDM's support of BLIMPs. This would be a substantial crowding in effect, equivalent in magnitude to crowding in half of India's non-CDM wind power capacity. This theoretical possibility does not undermine our finding that the CDM appears to have substantially misallocated resources. If anything, the possibility that marginal support could deliver indirect emissions reductions implies an even greater opportunity cost from subsidizing inframarginal projects.

Allowing for Mistakes.—One challenge to our findings is that, perhaps, we are applying our criteria too unsparingly. Maybe developers sometimes mistakenly invest in projects that are unlikely to be profitable, or the CDM mistakenly approves an application that they shouldn't have.

One way to address this concern is to change the definition of a BLIMP such that only subsidized projects that are *much* bigger and better located than the comparison projects qualify. An analysis capturing this adjustment was implemented in Figure 5.

Another method is to require that we see not just one but $N \geq 1$ inferior non-CDM projects before labeling something as a BLIMP. Row 19 of Table 2 shows the results for $N = 2$. The number of CDM projects that qualify as BLIMPs naturally falls when we demand to see two, three, four, or even five inferior wind farms. Even when setting $N = 5$, however, the BLIMPs that remain account for over 5 million CERs (see online Appendix Figure D4). A lottery assignment would still have a high probability of awarding fewer offsets to BLIMPs.

Partial Inframarginality.—The definition of a BLIMP already gives a great deal of deference to the CDM's assessment of which projects are marginal—we only question their assessment in cases where we can point to another wind farm, built in the same state and year with less generation capacity *and* lower capacity factor *and* greater connection distance. To give the CDM the greatest possible benefit of the doubt, we explore how our results are affected if we assume that every BLIMP is only partly inframarginal.

While we argue that BLIMPs would have been built regardless, they might not have been built with the same capacity. Had the CDM not promised an additional revenue stream, perhaps the developer would have ultimately failed to build the proposed 10 MW wind farm and instead been constrained to build only, say, an 8 MW project. This would be more likely in a context where there are diminishing returns to scale. Based on the understanding that in wind power generation there are increasing returns to scale, our main analysis counts all 10 MW as inframarginal. If this assumption is incorrect, it may be that only the 8 MW part of the project should be considered inframarginal.

To operationalize this idea, we use the next biggest inferior non-CDM project to bound from below the counterfactual capacity of BLIMPs. The next biggest non-CDM project represents the capacity that a BLIMP could have reached without CDM support. This approach yields a lower bound on inframarginal capacity.

Row 20 of Table 2 shows the results. Since we have built in the assumption that every BLIMP is only partly inframarginal, the blatantly inframarginal capacity is much lower—just under half of the number in our main results. The number of carbon credits allocated to support this inframarginal wind capacity is also lower—about a third of our main finding. This should be thought of as a conservative lower bound on a lower bound. Not only are BLIMPs a strict subset of inframarginal projects, but here we are also counting only the smallest possible fraction of each BLIMP as inframarginal. At this lower bound, a lottery mechanism would outperform the CDM about 30 percent of the time.

VI. Discussion

Our analysis reveals that the CDM program registered a large number of BLIMPs, larger even than would have been likely if registration was determined through a lottery mechanism. That we are able to come to these conclusions based on a small set of observable characteristics would seem to suggest that the regulator should also be well positioned to reject these same projects. At the very least, they should not perform worse than randomly selecting projects. This section considers their

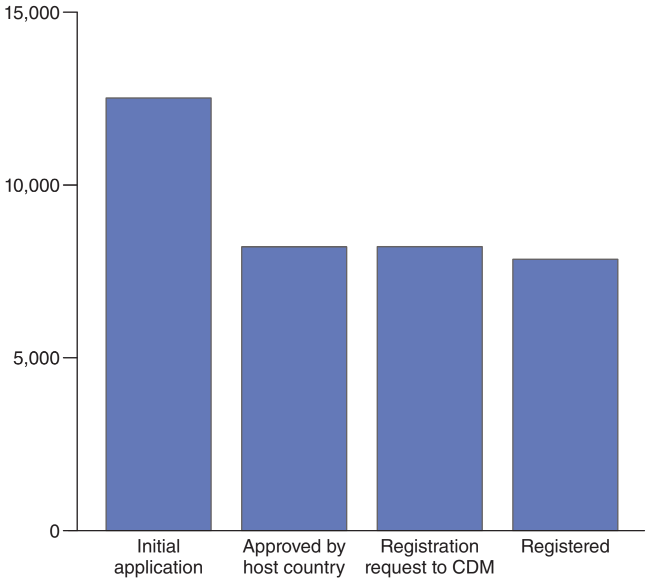

Figure 8. Number of CDM Applications at Each Stage

decision-making process with a view to understanding potential institutional failures that could explain our findings.

As discussed in Section IIA, projects must go through three stages of evaluation. First, the project has to be approved by the DNA. Its main task is to assess whether the project will help the host country achieve its sustainable development goals. Based on UNEP/Risoe's database, about 65 percent of applications across the world receive host country approval (Figure 8). After that, third party independent verifiers and the CDM Executive Board are meant to assess the "additionality" of the project (i.e., whether the project is marginal). In practice, verifiers sign off on and request registration for practically all host approved applications, and the CDM Executive Board approves over 95 percent of these requests.

Starting with the process of getting host country approval in India, diplomatic cables released by WikiLeaks in 2011 revealed that the Indian National Authority makes no independent assessment of whether a project is marginal or inframarginal and that the CDM Executive Board is aware of this.

> India grants host country approval to a CDM project based on the sustainability criteria following a presentation by the project developer to demonstrate that the project promotes economic, social, environmental and technological well-being . . . At a seminar on CDM in Mumbai, R K Sethi, Member Secretary of the National CDM Authority and the present Chairman of the CDM Executive Board, publicly admitted that the National CDM Authority takes the 'project developer at his word' for clearing the 'additionality' barriers (Consulate Mumbai 2008).

In practice, the structure of the screening process might actually favor inframarginal projects. The aforementioned diplomatic cables include evidence that project developers have had to secure credit before applying to the CDM (Consulate Mumbai 2008).

Perhaps none of this is a major concern for the CDM Executive Board because applications still have to be assessed by the third-party independent verifiers. These contractors, however, are selected and paid by the project developers, creating a strong incentive to approve or coach the projects into approval (Duflo et al. 2013). In practice, over 70 percent of applications are handled by just four companies—DNV, TÜV, BV Cert, and SGS. In the WikiLeaks report, we read that one verifier "urged companies to think of 'innovative' ways to qualify CDM projects" (Consulate Mumbai 2008). A member of DNV also reportedly "warned project developers to be more 'conservative' when estimating carbon credits that can be generated by the project, as projects which fall short of the projected number of credits approved by the Board can later be rejected" (Consulate Mumbai 2008).

Wara and Victor (2008) noted that "there is scant oversight on the integrity of the verification process and no record of punishing verifiers for misconduct." When the CDM Executive Board subsequently conducted spot checks on DNV, they discovered it could not provide evidence that qualified experts were involved in the validation work and that the basis for approval had not been documented (United Nations Framework Convention on Climate Change 2008). Similar deficiencies were later found at SGS and TÜV (Frunza 2013, 63). When an independent organization graded the major verifiers on a scale from A to F, TÜV received a D, SGS an E, and DNV and BV Cert an F (Schneider and Mohr 2009). As we have shown in Figure 8, these concerns extend well beyond the "big four" since verifiers have invariably recommended host-country-approved applications for CDM registration.

The final opportunity to separate inframarginal from marginal projects is when the ten-member CDM Executive Board considers each registration request. Figure 9 shows the number of registration requests made each year. Notice that the first request was submitted in September 2004, which means that projects listed as having registered before this date were approved retroactively, a sign that their completion was not contingent on CDM support. Between September 2004 and December 2013, the CDM Executive Board received 7,808 applications, requiring them to make approximately five decisions per day.[22] In practice, the largest volume of applications was submitted in the last year of the Kyoto Commitment Period, when the Board had to render decisions on nearly 15 applications per day. This is also when the average time from request to registration fell to its lowest level of just 17 days (Figure 9).

With no other sources of information, and limited time and resources to independently analyze complex financial and technical information, the Board was likely severely constrained in its ability to reject requests. In addition, board members are elected by the parties to the Kyoto Protocol, and it may be difficult for them to resist political pressures from host countries and offset buyers. Transparency International reported that "in closed-door meetings, board members have in some instances aggressively promoted projects that benefit their home countries or companies from their countries." (Transparency International 2011). For a sample of

[22] We assume a five-day work week accounting for 20 days of paid holiday and the ten paid holidays that the United Nations observes each year

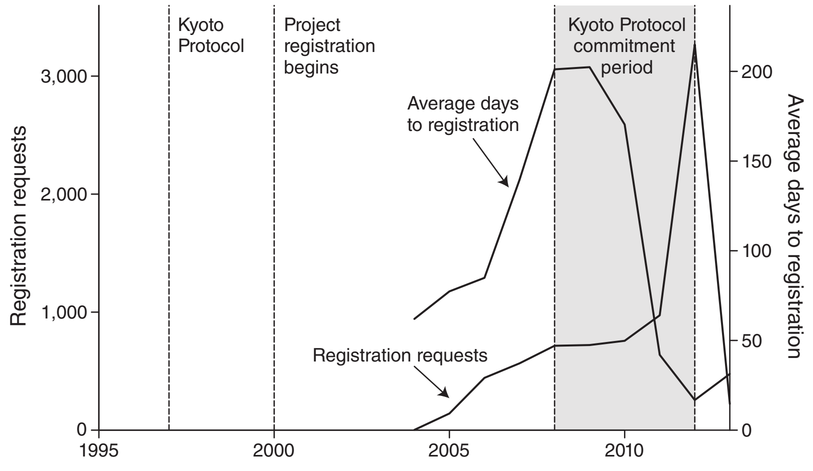

Figure 9. Registration Requests and Approval Times

1,000 early CDM projects, Flues, Michaelowa, and Michaelowa (2010) found that decisions systematically favor applications from board member countries, controlling for other proxies of project quality. Rather than maximizing additionality, the Board's de facto objective function appears to be one of maximizing the number of project approvals that it can credibly justify. Ultimately, over 95 percent of registration requests were approved.

The combined effect of these institutional failures is an almost automatic approval based on whether the host country approves the application in the first stage. In this context, the Indian DNA asks for no evidence of "additionality." The host country has a different objective function—maximizing contributions to its sustainable development goals. It is easy to see how these objectives can be at odds. Larger wind farms in windy locations may be inframarginal, but they also make a larger contribution toward providing affordable and sustainable energy.

The conflict between objective functions can be seen empirically by looking at the applications that didn't make it to the CDM Executive Board. We have been able to link 124 of these unsuccessful applications to Indian wind farms. The fact that these wind farms were built proves that they were inframarginal. Even so, only 43 percent of them would have been classified as BLIMPs had their applications been successful compared to our main finding of 54 percent of CDM projects.[23] The gap would be even larger if we considered the unsuccessful applications that were never built. The Indian DNA, then, systematically favored larger, better-located wind farms that were more likely to be BLIMPs. When combined with institutional failures at the level of the CDM Executive Board, this would account for our finding that the CDM has supported more BLIMPs than a lottery would have supported.

[23] To calculate this figure for unsuccessful applications, we first had to link unsuccessful applications to Indian wind farms, then change the CDM status of one project at a time and determine whether it would have been qualified as a BLIMP. Doing this exercise one project at a time means that the pool of non-CDM projects is almost unchanged from our main analysis, which makes the results more comparable.

VII. Conclusion

The last decade has seen billions of carbon offsets issued to project developers around the world, providing opportunities for regulatory compliance at a lower cost. However, when offset programs support projects that would have been developed anyway, they not only waste the limited resources available to mitigate climate change but also contribute to climate change by increasing global emissions. In the context of the CDM—the world's largest carbon offset program—we estimate that over half of approved carbon offsets for Indian wind farms were allocated to projects that would very likely have been built anyway. This is a substantial misallocation of resources.

When the CDM was created, India's wind power sector was identified as having huge potential for supporting marginal projects and increasing development beyond baseline projections. Yet we estimate that the CDM has approved at least 27 million tonnes worth of carbon offsets to inframarginal projects. If we extrapolate this rate of inframarginal support to the CDM as a whole, we calculate that the program may have allowed regulated polluters to emit an additional 6.1 billion tonnes of CO_2, equivalent to running 1,500 coal-fired power plants for a year. Given a social cost of carbon of \$190/tonne (EPA 2022; Rennert et al. 2022), the discounted economic losses associated with these emissions would be valued at \$1.2 trillion (\$2020).

We also find that the allocation of offsets to Indian wind power projects compares unfavorably with a lottery, indicating that there is substantial room for improvement in the design and implementation of the project selection mechanism. Accurately screening out projects that do not require subsidies is essential to safeguarding the integrity of offset programs.

Why is the CDM selecting these more profitable projects? Our analysis points to a series of institutional failures, starting with the Indian National Authority favoring larger wind projects in better locations in order to maximize contributions to its sustainable development goals and ending with a number of political and resource constraints that limit the CDM Executive Board's ability to independently vet and reject applications forwarded by host countries.

We acknowledge two important limitations of our analysis. First, it is possible that, even if awarding carbon offsets to these projects resulted in greater net emissions, the transfer of wealth from developed to developing countries may have increased welfare. Our analysis speaks only to whether or not the projects are likely to have been built without those offsets. Second, our analysis concerns only the direct effect of the CDM on emissions. We do not capture any indirect effects arising from support given to marginal projects that may have resulted in the establishment of credit and supply chains and enabling learning from others. Our analysis cannot rule out the existence of crowding in; however, our analysis implies that substantial crowding in would be required to compensate for the inframarginal projects that were supported by the CDM. This theoretical possibility does not undermine our finding that the CDM appears to have substantially misallocated resources. If anything, the possibility that marginal support could deliver indirect emissions reductions implies an even greater opportunity cost from subsidizing inframarginal projects.

REFERENCES

Agan, Amanda, and Sonja Starr. 2017. "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment." *Quarterly Journal of Economics* 133 (1): 191–235.

Anderson, Soren T., and James M. Sallee. 2011. "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards." *American Economic Review* 101 (4): 1375–1409.

Asher, Sam, and Paul Novosad. 2023. *Socioeconomic High-Resolution Rural-Urban Geographic Dataset for India (SHRUG) version 2.0*. Development Data Lab. https://www.devdatalab.org/shrug (accessed May 17, 2023).

Auffhammer, Maximilian, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel. 2013. "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change." *Review of Environmental Economics and Policy* 7 (2): 181–98.

Azoulay, Pierre, Joshua S. Graff Zivin, Danielle Li, and Bhaven N. Sampat. 2019. "Public R&D Investments and Private-Sector Patenting: Evidence from NIH Funding Rules." *Review of Economic Studies* 86 (1): 117–52.

Bayer, Patrick, Johannes Urpelainen, and Alice Xu. 2014. "Laissez Faire and the Clean Development Mechanism: Determinants of Project Implementation in Indian States, 2003–2011." *Clean Technologies and Environmental Policy* 16 (8): 1687–1701.

Bento, Antonio, Benjamin Ho, and Mario Ramirez-Basora. 2015. "Optimal Monitoring and Offset Prices in Voluntary Emissions Markets." *Resource and Energy Economics* 41: 202–23.

Bento, Antonio M., Ravi Kanbur, and Benajmin Leard. 2015. "Designing Efficient Markets for Carbon Offsets with Distributional Constraints." *Journal of Environmental Economics and Management* 70: 51–71.

Berry, Steven T. 1992. "Estimation of a Model of Entry in the Airline Industry." *Econometrica* 60 (4): 889–917.

Bharadwaj, Prashant, Leah K. Lakdawala, and Nicholas Li. 2019. "Perverse Consequences of Well Intentioned Regulation: Evidence from India's Child Labor Ban." *Journal of the European Economic Association* 18 (3): 1158–95.

Black, Richard, Kate Cullen, Byron Fay, Thomas Hale, John Lang, Saba Mahmood, and Steve M. Smith. 2021. *Taking Stock: A Global Assessment of Net Zero Targets*. Oxford, UK: University of Oxford.

Blanco, María Isabel. 2009. "The Economics of Wind Energy." *Renewable and Sustainable Energy Reviews* 13 (6–7): 1372–82.

Bloom, Nick, Rachel Griffith, and John Van Reenen. 2002. "Do R&D Tax Credits Work? Evidence from a Panel of Countries 1979–1997." *Journal of Public Economics* 85 (1): 1–31.

BNEF. 2013. *Bloomberg New Energy Finance, Renewable Energy Project Database*. New York, NY: BNEF. https://about.bnef.com/ (accessed July 2013).

Boomhower, Judson, and Lucas W. Davis. 2014. "A Credible Approach for Measuring Inframarginal Participation in Energy Efficiency Programs." *Journal of Public Economics* 113: 67–79.

Burlig, Fiona, Akshaya Jha, and Louis Preonas. 2020. *Data and Code for: "Transmission Constraints and Electricity Trade in India."* Chicago, IL: EPIC. https://epic.uchicago.in/project/transmission-constraints-and-electricity-trade-in-india/ (accessed October 8, 2019)

Bushnell, James B. 2010. "The Economics of Carbon Offsets." NBER Working Paper 16305.

Calel, Raphael, Jonathan Colmer, Antoine Dechezleprêtre, and Matthieu Glachant. 2025. *Data and Code for: "Do Carbon Offsets Offset Carbon?"* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E195621V1.

Carley, Sanya, Lincoln L. Davies, David B. Spence, and Nikolaos Zirogiannis. 2018. "Empirical Evaluation of the Stringency and Design of Renewable Portfolio Standards." *Nature Energy* 3 (9): 754–63.

Chadwick, Bruce P. 2006. "Transaction Costs and the Clean Development Mechanism." *Natural Resources Forum* 30 (4): 256–71.

Chandra, Ambarish, Sumeet Gulati, and Milind Kandlikar. 2010. "Green Drivers or Free Riders? An Analysis of Tax Rebates for Hybrid Vehicles." *Journal of Environmental Economics and Management* 60 (2): 78–93.

Consulate Mumbai. 2008. "Carbon Credits Sufficient but Not Necessary for Sustaining Clean Energy Projects of Major Indian Business Groups." Wikileaks ID: 08MUMBAI340. https://wikileaks.org/plusd/cables/08MUMBAI340_a.html.

Cullen, Joseph. 2013. "Measuring the Environmental Benefits of Wind-Generated Electricity." *American Economic Journal: Economic Policy* 5 (4): 107–33.

Davis, Lucas W. 2008. "The Effect of Driving Restrictions on Air Quality in Mexico City." *Journal of Political Economy* 116 (1): 38–81.

Dechezlepretre, Antoine, Elias Einio, Ralf Martin, Kieu-Trang T. Nguyen, and John Van Reenen. 2016. "Do Tax Incentives for Research Increase Firm Innovation? An RD Design for R&D." NBER Working Paper 22405.

Doleac, Jennifer, and Benjamin Hansen. 2020. "The Unintended Consequences of 'Ban the Box': Statistical Discrimination and Employment Outcomes When Criminal Histories Are Hidden." *Journal of Labor Economics* 38 (2): 321–74.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2013. "Truth-Telling by Third-Party Auditors and the Response of Polluting Firms: Experimental Evidence from India." *Quarterly Journal of Economics* 128 (4): 1499–1545.

EPA. 2022. *Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances*. Washington, DC: EPA.

Fearnehough, Harry, Lambert Schneider, and Carsten Warnecke. 2021. *The Potential Impact of Transitioning CDM Units and Activities to the Paris Agreement* (*webinar*). Freiburg im Breisgau, Germany: Öko-Institut.

Filmer, Deon, Jed Friedman, Eeshani Kandpal, and Junko Onishi. 2023. "Cash Transfers, Food Prices, and Nutrition Impacts on Ineligible Children." *Review of Economics and Statistics* 105 (2): 327–43.

Fischer, Carolyn. 2005. "Project-Based Mechanisms for Emissions Reductions: Balancing Trade-Offs with Baselines." *Energy Policy* 33 (14): 1807–23.

Flues, Florens, Axel Michaelowa, and Katharina Michaelowa. 2010. "What Determines UN Approval of Greenhouse Gas Emission Reduction Projects in Developing Countries? An Analysis of Decision Making on the CDM Executive Board." *Public Choice* 145 (1/2): 1–24.

Frunza, Marius-Christian. 2013. *Fraud and Carbon Markets: The Carbon Connection*. London, UK: Routledge.

GIZ. 2014. *Data for: "Carbon Market Roadmap for India: Looking back on CDM and Looking ahead."* https://doi.org/10.3886/E195621V1.

Gowrisankaran, Gautam, Stanley S. Reynolds, and Mario Samano. 2016. "Intermittency and the Value of Renewable Energy." *Journal of Political Economy* 124 (4): 1187–1234.

Hall, Bronwyn, and John Van Reenen. 2000. "How Effective are Fiscal Incentives for R&D? A Review of the Evidence." *Research Policy* 29 (4–5): 449–69.

Holland, Stephen P., Erin T. Mansur, Nicholas. Z. Muller, and Andrew J. Yates. 2016. "Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors." *American Economic Review* 106 (12): 3700–29.

Howell, Sabrina T. 2017. "Financing Innovation: Evidence from R&D Grants." *American Economic Review* 107 (4): 1136–64.

Ito, Koichiro, and James M. Sallee. 2018. "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel Economy Standards." *Review of Economics and Statistics* 100 (2): 319–36.

Jack, B. Kelsey, and Seema Jayachandran. 2019. "Self-Selection into Payments for Ecosystem Services Programs." *PNAS* 116 (12): 5326–33.

Jack, B. Kelsey, Seema Jayachandran, Namrata Kala, and Rohini Pande. Forthcoming. "Money (Not) to Burn: Payments for Ecosystem Services to Reduce Crop Residue Burning." *American Economic Review: Insights*.

Jagadeesh, A. 2000. "Wind Energy Development in Tamil Nadu and Andhra Pradesh, India Institutional Dynamics and Barriers – A Case Study." *Energy Policy* 28 (3): 157–68.

Jayachandran, Seema. 2013. "Liquidity Constraints and Deforestation: The Limitations of Payments for Ecosystem Services." *American Economic Review* 103 (3): 309–13.

Jayachandran, Seema, Joost De Laat, Eric F. Lambin, Charlotte Y. Stanton, Robin Audy, and Nancy E. Thomas. 2017. "Cash for Carbon: A Randomized Trial of Payments for Ecosystem Services to Reduce Deforestation." *Science* 357 (6348): 267–73.

Kathuria, Vinish, Pradeep Ray, and Rekha Bhangaonkar. 2015. *Data for: "FDI (Foreign Direct Investment) in Wind Energy Sector in India: Testing the Effectiveness of State Policies Using Panel Data." Energy* 80: 190–202 (accessed March 19, 2021).

Kossoy, Alexandre, Grzegorz Peszko, Klaus Oppermann, Nicolai Prytz, Noemie Klein, Kornelis Blok, Long Lam, and et al. 2015. *State and Trends of Carbon Pricing 2015*. Washington, DC: World Bank.

Kwoka, John E. 1983. "The Limits of Market-Oriented Regulatory Techniques: The Case of Automotive Fuel Economy." *Quarterly Journal of Economics* 98 (4): 695–704.

Lee, Joyce, Feng Zhao, Alastair Dutton, Ben Backwell, Ramon Fiestas, Liming Qiao, Naveen Balachandran, and et al. 2021. *Global Wind Report 2021*. Brussels, Belgium: Global Wind Energy Council.

Mast, Evan. 2020. "Race to the Bottom? Local Tax Break Competition and Business Location." *American Economic Journal: Applied Economics* 12 (1): 288–317.

Michaelowa, Axel, and Pallav Purohit. 2007. "Additionality Determination of Indian CDM Projects: Can Indian CDM Project Developers Outwit the CDM Executive Board?" Climate Strategies Final Paper.

Michaelowa, Axel, Igor Shishlov, and Dario Brescia. 2019. "Evolution of International Carbon Markets: Lessons for the Paris Agreement." *WIREs Climate Change* 10 (6): e613.

Ministry of Power, Government of India. 2013. *Data from: "State Distribution Utilities First Annual Integrated Rating."* Technical Report. https://doi.org/10.3886/E195621V1.

Montero, Juan-Pablo. 2000. "Optimal Design of a Phase-in Emissions Trading Program." *Journal of Public Economics* 75 (2): 273–91.

Moretti, Enrico, and Daniel J. Wilson. 2014. "State Incentives for Innovation, Star Scientists, and Jobs: Evidence from Biotech." *Journal of Urban Economics* 79: 20–38.

Muñoz-Sabater, J., E. Dutra, A. Agustí-Panareda, C. Albergel, G. Arduini, G. Balsamo, S. Boussetta, and et al. 2019. *ERA5-Land Hourly Data from 1950 to Present*. Reading, UK: Copernicus Knowledge Base. https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-land?tab=overview (accessed March 8, 2021).

Nunn, Nathan, and Diego Puga. 2012. *Data for: "Ruggedness: The Blessing of Bad Geography in Africa." Review of Economics and Statistics* 94 (1): 20–36 (accessed May 17, 2023).

Oliva, Paulina. 2015. "Environmental Regulations and Corruption: Automobile Emissions in Mexico City." *Journal of Political Economy* 123 (3): 686–724.

Parker, Dominic P., and Bryan Vadheim. 2017. "Resource Cursed or Policy Cursed? US Regulation of Conflict Minerals and Violence in the Congo." *Journal of the Association of Environmental and Resource Economists* 4 (1): 1–49.

Phadke, Amol, Ranjit Bharvirkar and Jagmeet Khangura. 2011. *Reassessing Wind Potential Estimates for India: Economic and Policy Implications*. Lawrence Berkeley National Laboratory.

Pless, Jacquelyn. 2021. "Are 'Complementary Policies' Substitutes? Evidence from R&D Subsidies in the UK." Unpublished.

Point Carbon. 2010. Carbon 2010 - Return of the Sovereign. http://sa.indiaenvironmentportal.org.in/files/Carbon%202010.pdf.

Purohit, Pallav, and Axel Michaelowa. 2007. "Potential of Wind Power Projects under the Clean Development Mechanism in India." *Carbon Balance and Management* 2 (1): 8.

Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, and et al. 2022. "Comprehensive Evidence Implies a Higher Social Cost of CO2." *Nature* 610: 687–92.

Rosenbaum, Paul R., and Jeffrey H. Silber. 2009. "Amplification of Sensitivity Analysis in Matched Observational Studies." *Journal of the American Statistical Association* 104 (488): 1398–1405.

Ryan, Nicholas. 2022. "Holding Up Green Energy: Counterparty Risk in the Indian Solar Power Market." Unpublished.

Schneider, Lambert. 2009. "Assessing the Additionality of CDM Projects: Practical Experiences and Lessons Learned." *Climate Policy* 9 (3): 242–54.

Schneider, Lambert Richard. 2011. "Perverse Incentives under the CDM: An Evaluation of HFC-23 Destruction Projects." *Climate Policy* 11 (2): 851–64.

Schneider, Lambert, and Anja Kollmuss. 2015. "Perverse Effects of Carbon Markets on HFC-23 and SF6 Abatement Projects in Russia." *Nature Climate Change* 5 (12): 1061–63.

Schneider, Lambert, and Lennart Mohr. 2009. *A Rating of Designated Operational Entities* (*DOEs*) *Accredited under the Clean Development Mechanism* (*CDM*)*: Scope, Methodology and Results*. Freiburg im Breisgau, Germany: Öko-Institut.

Shapiro, Joseph S., and Reed Walker. 2020. "Is Air Pollution Regulation Too Stringent?" Unpublished.

Slattery, Cailin. 2019. "Bidding for Firms: Subsidy Competition in the US." Unpublished.

Slattery, Cailin, and Owen Zidar. 2020. "Evaluating State and Local Business Incentives." *Journal of Economic Perspectives* 34 (2): 90–118.

Taylor, Rebecca L. C. 2019. "Bag Leakage: The Effect of Disposable Carryout Bag Regulations on Unregulated Bags." *Journal of Environmental Economics and Management* 93: 254–71.

Transparency International. 2011. *Global Corruption Report: Climate Change*. London, UK: Earthscan.

UNEP DTU. 2021. *CDM Pipeline Database*. Copenhagen, Denmark: UNEP DTU. https://unepccc.org/cdm-ji-pipeline/ (accessed May 2021).

United Nations Framework Convention on Climate Change. 2008. *Minutes of the Executive Board of the Clean Development Mechanism Meeting 44 – Annex 2: List of Non-conformities of DNV*. Bonn, Germany: United Nations Framework Convention on Climate Change. https://cdm.unfccc.int/EB/archives/meetings_08.html (accessed November 28, 2023).

Van Benthem, Arthur, and Suzi Kerr. 2013. "Scale and Transfers in International Emissions Offset Programs." *Journal of Public Economics* 107: 31–46.

Wara, Michael. 2007a. "Is the Global Carbon Market Working?" *Nature* 445 (7128): 595–96.

Wara, Michael. 2007b. "Measuring the Clean Development Mechanism's Performance and Potential." *UCLA Law Review* 55: 1759–1803.

Wara, Michael, and David G. Victor. 2008. "A Realistic Policy on International Carbon Offsets." Unpublished.

World Bank Group. 2019. *Data for: "State and Trends of Carbon Pricing 2019."* World Bank, Washington, DC. https://documents1.worldbank.org/curated/en/191801559846379845/pdf/State-and-Trends-of-Carbon-Pricing-2019.pdf (accessed August 14, 2021).