

The Economic Impacts of Clean Energy*

Costas Arkolakis Conor Walsh

November 9, 2024

Abstract

In this paper we assess the economic impacts of moving to a renewable-dominated grid in the US. We use projections of capital costs to develop price bounds on future wholesale power prices at the local geographic level. We then use a class of spatial general equilibrium models to estimate the effect on wages and output of prices falling below these bounds in the medium term. Power prices fall anywhere between 20% and 80%, depending on local solar resources, leading to an aggregate real wage gain of 2-3%. Over the longer term, we show how moving to clean energy represents a qualitative change in the aggregate growth process, alleviating the “resource drag” that has slowed recent productivity growth in the US.

Keywords: Renewable Energy, Solar Energy, Sufficient Statistics, Economic Growth

*We thank our editors Jón Steinsson and Janice Eberly, as well as our discussants Neil Mehrotra and Koichiro Ito for their excellent comments. Cheolhwan Kim, Pablo Mones, Eshaan Patel, and Yirui Zhao provided outstanding research assistance. We would also like to thank Lucas Epstein for various comments. This conference draft was presented at the Fall 2024 Brookings’ Papers for Economic Activity (BPEA) Conference, and the final version will be published in the Fall 2024 BPEA issue. The authors would like to acknowledge the support of NSF under grant SES-2117288. Other than the aforementioned, the authors did not receive financial support from any firm or person for this article or from any firm or person with a financial or political interest in this article. The authors are not currently an officer, director, or board member of any organization with a financial or political interest in this article.

1 Introduction

The US electricity grid is undergoing a historic transition. The share of renewable electricity in generation has begun to rise rapidly from virtually nothing a decade ago (see Figure 1), supported both by policy and significant falls in the capital costs of solar and wind. This trend has been widely celebrated for the climate benefits it brings with it; when solar and wind displace coal and gas, CO₂ emissions fall.

However, both the academic literature and popular analysis have placed somewhat less emphasis on the economic impacts of the transition to a clean grid. Chief among these are lower wholesale power prices, and a displacement of fossil fuel activity. The former is likely to lead to greater economic production, higher wages, and cheaper goods prices. The latter may cause displacement of fossil fuel employment, and transitional pressures as these workers retrain and shift into other sectors.

In this paper, we assess the first impact. We begin by developing projections of all-inclusive capital costs for firmed solar, i.e. solar backed by storage, in the near future. We then use these projections to construct bounds on wholesale electricity prices across the US at a relatively fine geographic scale. We show that a move to a grid dominated by firmed solar power is likely to see substantially lower wholesale electricity prices in most areas of the United States, with power costs falling between 20% and 80% depending on solar insolation and local land costs.

In a second step, we use these price bounds in a class of general equilibrium models to estimate the impact on local wages and production of moving to a solar-dominated grid. In the medium term (out to 2040), we find a fairly substantial increase in wages (on the order of 2-3% nationwide), with large regional heterogeneity. Rural areas stand to benefit the most, owing to a greater share of electricity in the factor inputs of their industry concentration mixes. However, many large cities and counties in Texas and California also see average wage rises of almost 5%, owing to substantial decreases in wholesale prices in these sunny areas.

Finally, we consider the impact in the longer term. We outline a conceptual feature of the coming era of clean growth that is qualitatively different to the recent era of fossil-fueled growth. Namely, renewable capital accumulation relieves the aggregate drag of finite resource extraction and rising energy prices. A significant part of innovative effort since the 1970s has been directed at increasing energy efficiency to offset rising prices. When electricity comes from zero marginal cost sources, such as sunshine and wind, rising resource prices stop constraining growth. The economy's innovative resources then

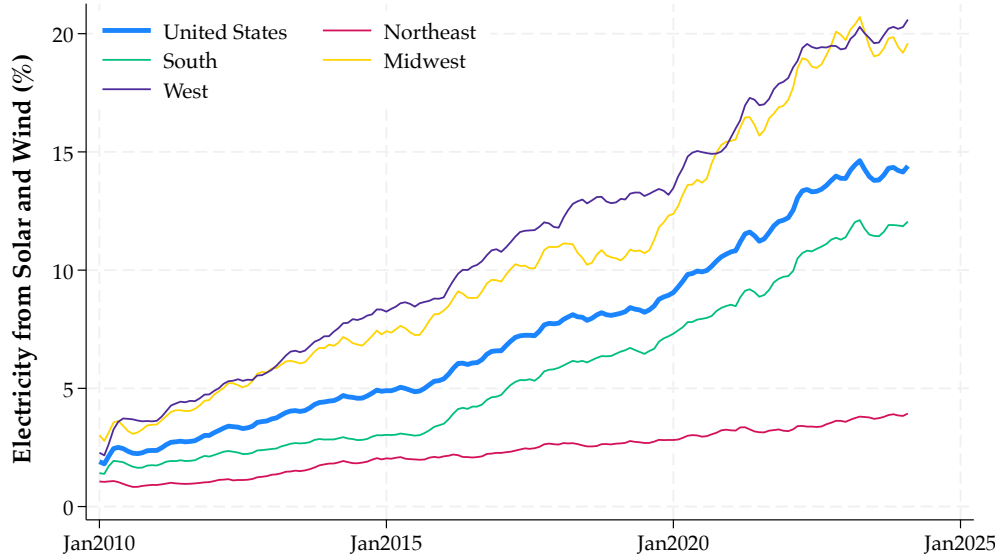
redirect away from energy-specific technical change to more general progress, raising the aggregate growth rate.

In making the points of this paper, we are required to make some assumptions about the future path of technology. While renewable energy is already the cheapest source of bulk energy supply at many points on the US grid, much depends on what happens to capital costs from this point forward, and in particular on continued falls in the cost of storage. We make some assumptions that we believe are reasonable, but there will be many points over which reasonable people can disagree, and we don't resile from this. Our purpose in this paper is to ask "What if?", and think about a world in which renewable energy is cheap, abundant and the dominant source of power in the US. We take recent technological trends in energy seriously, and analyze their impact through the lens of economic theory. The landscape is littered with the bones of prognosticators who wrote off renewable progress, and we hope not to join that graveyard.

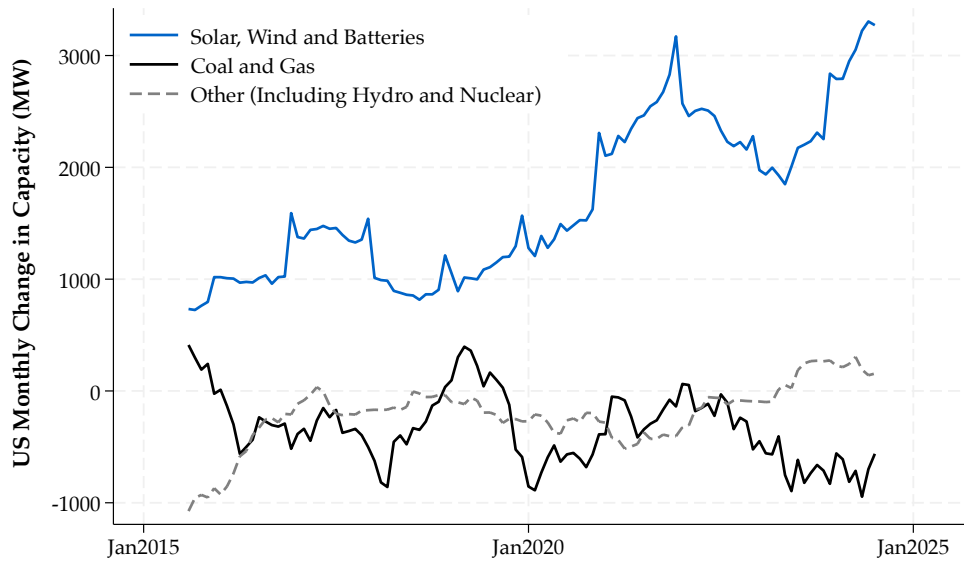
Related Literature. Our approach to evaluate the economic implications of the energy transition differs significantly from other approaches in the literature, most notably Jenkins et al. (2022), Bistline et al. (2023a), Bistline et al. (2023b) and Abhyankar et al. (2021). These papers use detailed engineering and energy system models to compute the implications of economic stimulus policies and renewable subsidies on energy prices and renewable uptake. They build analyses of supply curves and transmission from the ground up, along with modeling the use of energy in production, and study least cost investment approach pathways to achieving net-zero under various technological assumptions and policy scenarios. The literature around energy systems models is vast and influential; see Pfenninger et al. (2014) for a review. Recent applications of to this approach to specifically study the labor market impacts of renewable penetration include Jenkins et al. (2021) and Mayfield et al. (2023). Complementary to this approach is empirical work by Hanson (2023), who measures the local labor market effects of initial exposure to coal production.

Instead, we use projections of firmed solar capital costs in the near future to develop spatial bounds on future wholesale prices, and then incorporate them into a general equilibrium spatial model. In this sense, our approach is most closely related to our current work in Arkolakis and Walsh (2023). However, instead of developing a model of the grid and transmission of energy across space, we focus on the local bounds as a measure of energy cost changes across space. We then develop sufficient statistics to trace the impact of electricity price shocks onto wages at fine levels of disaggregation, without actually having to estimate and solve a fully-specified economic model. We view this as complementary to the successful energy systems modeling approach above, as while we abstract from the

Figure 1: The Renewable Transition in the US



(a) Solar and Wind Share in Electricity Generation



(b) Total Change In Capacity By Technology

Notes: Panel (a) of this figure shows the 12-month moving average of the monthly share of total electricity at the regional level coming from solar and wind. Panel (b) shows the 12-month moving average of the change in nameplate capacity by technology. Data are from the US Energy Information Administration.

detail and computational rigor imposed by these models, our simpler approach has the advantage of being readily interpretable, and can act as a basic starting point to shape analysis and policy.

There is an important economic literature that builds more aggregate macroeconomic models to study the energy transition. Integrated assessment models have long studied the interaction between fossil fuel extraction and the macroeconomy, beginning with the DICE model of Nordhaus (1993), and updated most recently in Barrage and Nordhaus (2024).¹ Work by Desmet et al. (2018), Bilal and Rossi-Hansberg (2023) and Cruz and Rossi-Hansberg (2024) has pushed this literature to disaggregate the effects of climate change and production shifts into heterogeneous spatial impacts within countries, while maintaining the discipline imposed by general equilibrium in the aggregate.² In addition, recent work by Mehrotra (2024) uses updated technology cost assumptions to show that the macroeconomic costs of transitioning to a net-zero economy are far smaller than supposed even recently. Our work attempts to add to both of these research strands.

A separate literature endogenizes the direction of technological change in energy innovation, building off the foundational theory of Acemoglu (2002), and applied to the context of energy and the environment in Acemoglu et al. (2012) (see Gillingham et al. (2008) for a review of earlier models). This literature continues to be active, with other important papers in this direction being Lemoine (2024), Kanzig and Williamson (2024) and Acemoglu et al. (2023). An influential paper by Hassler et al. (2021) models the slowdown in broader innovation and the increase in energy-saving technical change since the energy price shocks of the 1970s. We build on this modeling framework in the final section and show that when energy is provided by accumulable capital, as in the case of renewables, rather than exhaustible fossil fuels, innovation resources can be redirected to broad innovation and increase the aggregate growth rate.

2 Wholesale Power Prices In the Medium Term

Renewable energy has begun to grow rapidly in the US. Figure 1 shows that the share of electricity produced by wind and solar has risen in all regions in the US, from virtually nothing in 2010 to around 14% in 2024, with several regions seeing much higher penetration. In the bottom panel we show the aggregate flows of investment into renewables and

¹A strand of the literature emphasizes the endogenous effects that access to electricity has on output through energy prices, and incorporates supply and demand of energy in a macroeconomic model (Nordhaus et al., 1973; Kypreos and Bahn, 2003; Edenhofer et al., 2013).

²See Desmet and Rossi-Hansberg (2024) for a recent review.

non-renewables. The increase in electricity generation is being driven by rapid increases in nameplate capacity for solar, wind and batteries. Fossil fuels as a group are shrinking, while hydro and nuclear remain roughly stable. In Figure 17 in the Appendix we break these flows out further. Within fossil fuels, coal has been rapidly exiting the grid, while gas capacity continued to be built up until 2022 at roughly the same speed as wind energy, though in the last year new investment has dropped off rapidly (see also Figure 19). Solar investment has been consistently increasing over the last 10 years, to the point that in 2024 it currently accounts for the most new capacity investment out of all power technologies. Utility-scale batteries are a very recent addition to the grid, but in the last two years have begun to scale up rapidly, supporting the intermittent energy flows of solar and wind.

Analyzing the long-run price impact of moving to a renewable-dominated grid is a challenging endeavor. The US electric grid is enormously complicated, consisting of many interlocking organizations and systems. The transition to renewables will have heterogeneous impacts on prices in different locations, depending on local renewable resources, the pricing mechanisms of local utilities, and the strength of local transmission networks (a point we explored in earlier work in Arkolakis and Walsh (2023)).

Here, we try and cut through some of this complexity. We consider a simple bounds approach that is helpful as a guide to shape thinking. This approach abstracts from most of the grid’s complexity, and starts from the observation that capital costs are the dominant direct cost in supplying renewable energy.³ Operating costs of renewables are negligible, and depreciation and maintenance expenses have proved to be very low. As a result, upfront capital costs determine the economics of supply. In deregulated markets, particularly those that have implemented Locational Marginal Pricing, such costs place an upper bound on future steady state wholesale prices at any point on the grid.⁴

To make this point, we first note that renewable power generation capital is unlike conventional fossil fuel generation assets in several respects. Three stand out in particular, and form the basis for our analysis in this section.

First, renewables are modular. By this we mean that the generating unit comes in small sizes available at constant fixed prices, many of which are strung together to form a plant.

³Throughout this paper we use renewable energy to refer solely to photovoltaic (PV) solar power and onshore wind. These technologies are widely considered to be the dominant technologies in the medium term, with offshore wind playing a more limited role in the US due to geography and regulatory constraints. Other renewable technologies are either early stage and not cost competitive with PV and onshore wind (such as geothermal and wave energy), or like hydroelectricity are mature with limited scope for expansion.

⁴We abstract from local distribution costs and fixed network charges which show up in retail prices for this analysis, and return to them briefly in Section 3.5.

In contrast, fossil fuel assets such as coal-fired power stations tend to be large, complex installations with substantial fixed costs. This historically lead to a structure of centralized generation in large plants, with transmission lines strung to load centers. The modularity of renewables makes it easier to build them in smaller sized plants, and facilitates a much more decentralized grid. While large installations certainly exist, recent renewable projects tend to be of varying small and medium sizes, and are more dispersed around load areas.⁵

Second, fuel costs are zero, and the productivity of the asset depends mainly on where it is placed in space. Placing a solar panel in the sunshine or a wind turbine on a gusty ridge occasions zero direct input costs over the life of the asset. However, electricity output will differ widely across the country. The productivity of a solar panel in terms of total annual electricity production is around two times higher in Arizona than in Maine (see Figure 18 in the Appendix). The divergence in wind potential across space is even starker. Average wind power output is a cubic in average wind speeds. As a result, a wind turbine in the windiest locations, such as South Dakota, will produce around 5 times the electricity of the least windy locations, such as Florida.

Third, renewables are intermittent. As renewable penetration increases, more backup from storage or rapid-response peaking plants is required. In what follows, we will assume that in the medium run, renewables are completely backed up by battery storage, and examine the cost implications.⁶

Combining these assumptions allows us to develop a simple asset pricing equation that must hold in the long run wherever renewables are installed on the grid. Let $Q_{\ell t}$ denote the all-in capital cost of a megawatt (MW) of firmed renewable capital in location ℓ , whether solar or wind, inclusive of storage costs. Let θ_ℓ be the expected annual output of the capital unit in megawatt-hours (MWh), and $p_{\ell,t}^\mathcal{E}$ the average price of a MWh of electricity in location ℓ at time t . Assuming there are annual depreciation costs that occur at rate δ , and a cost of financing $R_{0 \rightarrow t}$ (where this should be read as the cumulative compound interest rate from year 0 until year t , namely $R_{0 \rightarrow t} \equiv \prod_{\tau=1}^t (1 + r_\tau)$), we can write.

$$Q_{\ell t} = \sum_{\tau=1}^T R_{t \rightarrow t+\tau}^{-1} (1 - \delta)^\tau \theta_\ell p_{\ell,t+\tau}^\mathcal{E}, \quad (1)$$

⁵We will also assume that their modularity by its very nature encourages a reduction in the market power of large incumbents. Bahn et al. (2021) caution that if renewable investment is developed primarily by legacy fossil fuel incumbents, the effect on wholesale prices of lower generation costs could be muted.

⁶Around half of the current solar projects in the interconnection queue are hybrid plants with a storage component, up from none just five years ago; see the Lawrence Berkeley National Laboratory here.

where T is the lifespan of the project (typically around 30 years for solar panels, and somewhat longer for wind energy). Appropriate adjustments can be made to incorporate longer times to build; typically, once approved a new solar plant can be constructed in 12-24 months, with a wind project taking somewhat longer.

Crucially, this equation does not have to hold everywhere on the grid in the medium term. There will be many places where wholesale power costs are lower than what would be implied by local capital costs, particularly in dense urban regions and places with poor renewable resources. In that case, equation (1) would be an inequality, with the left-hand side being greater than the right. In such places, electricity will be imported from other low-cost regions, with the ability to access their low prices driven crucially by transmission capacity. In addition, equation (1) may hold for solar in sunny regions, and wind in windy regions, but it does not necessarily have to hold for both at the same time.

In places where it does hold, we can solve for the steady price of power that must occur in these regions, by setting the wholesale electricity price $p_{\ell,t+s}^{\mathcal{E}} = p_{\ell}^{\mathcal{E}}$ to its medium-run average, and then writing

$$p_{\ell}^{\mathcal{E}} = \bar{Q}_{\ell} \left(\theta_{\ell} \sum_{\tau=1}^T R_{0 \rightarrow \tau}^{-1} (1 - \delta)^{\tau} \right)^{-1}, \quad (2)$$

where we let \bar{Q}_{ℓ} be the upfront medium-run investment cost of renewable capital, backed up by storage to address its inherent intermittency.⁷ This shows that in the long run, in places where renewables are installed, there are two essential determinants of electricity prices across space: upfront capital costs \bar{Q}_{ℓ} , and potential expected annual output, θ_{ℓ} .

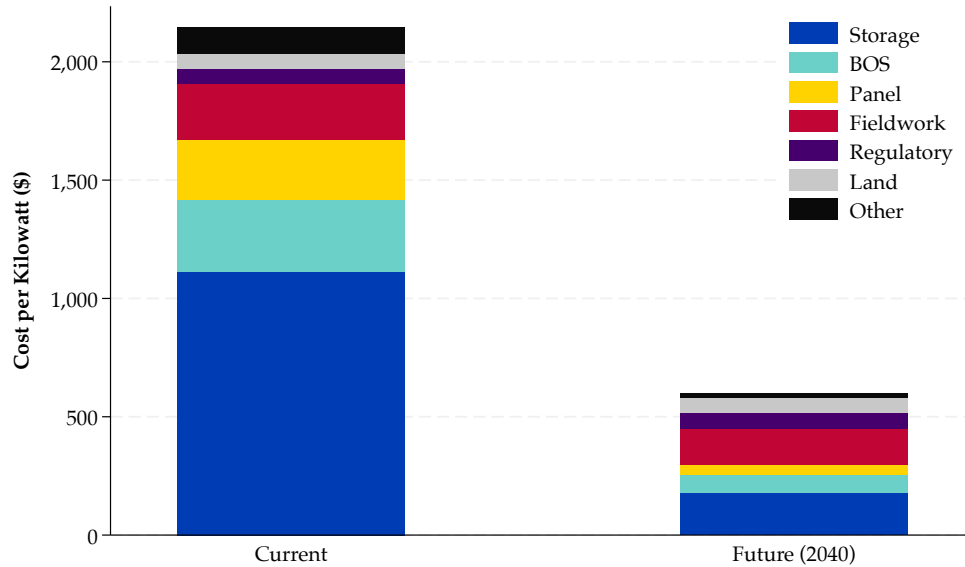
Leaving aside potential for the moment, one can think about breaking down the capital cost into several components that we can size across different areas:

$$\begin{aligned} \bar{Q}_{\ell} = & \text{Plant Capital Cost} + \text{Balance of System Cost} \\ & + \text{Construction Cost} + \text{Land Cost} \\ & + \text{Regulatory Cost} + \text{Storage Cost.} \end{aligned} \quad (3)$$

We can try to get a handle on each element of these, and think how they might change into the future. For the rest of the paper, we are going to focus on solar power. While

⁷This formula is very similar to what is often called the levelized cost of energy (LCOE) in the literature, with the main difference being how depreciation expenses are treated.

Figure 2: Firmed Solar Project Costs Now and in the Future



Notes: Figure shows the breakdown of upfront installation costs of a kilowatt of solar power. Current uses data from NREL in Ramasamy et al. (2023), along with author calculations. Future is based on author calculations. Prices are in 2023 dollars.

we expect wind power to be an important part of the generation mix of the future, the recent explosion of solar power and its pairing with lithium-ion storage leads us to believe that it will be the dominant technology. Its capital costs have been falling faster than the capital costs of wind for a prolonged period, and in many places it is now the cheapest form of unsubsidized bulk energy supply. As can be seen in Figure 19, the committed and under-construction project pipeline in the US is dominated by solar projects and short-duration lithium-ion storage, with wind a distant second (and fossil investment being virtually absent). Nonetheless, the techniques we use here can easily be adapted to study the impact of wind energy, and we do so in Appendix C.

In Figure 2 we show the current breakdown of an installed unit of firmed solar capital at utility scale. Our source for this is the National Renewable Energy Laboratory (NREL) “U.S. Solar Photovoltaic System and Energy Storage Cost Benchmarks” for 2023. We make three small adjustments to the estimate of this paper. First, we use the price for solar panels coming from data from Bloomberg New Energy Finance in 2023, reflecting significant falls in 2023 relative to 2022 (the data used by NREL). Second, while NREL includes land leasing costs in operation and maintenance, we cumulate them and discount them to include them in upfront costs. We separately include depreciation and maintenance in δ on the right hand side of equation (2). Last, we extend the amount of storage for each unit of solar from 2.5 hrs to 8 hours, using costs from Bloomberg New Energy Finance in 2023.

Allowing for 8 hours of storage per unit of solar capital effectively makes each hour of sunlight captured in any day across the year completely dispatchable on demand. The average daily output of a 1 kilowatt (kW) solar panel across the United States is 4 kilowatt-hours (kWh), with substantial heterogeneity across the country (see Figure (18)). This rises to around 6 kWh in the summer months, and falls to 2 kWh in the middle of winter. With 8 hours of storage for a 1kW panel up to 8kWh of output can be stored, so that solar power can provide for round-the-clock power (some panels supply while the sun is up, others dispatch their stored output in the evening or in the early morning hours), with a buffer for intra-week variability.


It is worth stressing this, since there is often some confusion in popular discussion of this point. Fully dispatchable solar power does not require 24 hours of storage for each solar panel. All it requires is enough storage so that each hour of sunlight captured in a day can be dispatched at will. In a crude example, having 5kWh of usable sunlight a day means that 5kW of solar panels, each with 5 hours of storage, can supply 1kW of power continuously throughout the day. Increasing the storage buffer to 8 hours, or building an extra 3kW of panels with 5 hours of storage, provides a reserve for cloudier days.

What about seasonal variability? The same location that has an average 5kWh of sunlight throughout the year might see 3kWh in winter and 7kWh in summer. Lithium-ion batteries are not ideal for long-term storage to offset this seasonable variation, as efficient use of the asset requires continual charging and discharging cycles. To the extent that gas backup is less available in the winter in a future renewable-dominated grid, this issue can be dealt with via the combination of *overbuilding* and *curtailment*. This involves building enough firmed solar to meet winter demand levels, and then in summer curtailing (or shutting off) the excess generation.⁸

In practice, for our exercise what this means is lowering the “capacity factors” of firmed solar. A capacity factor tells us what fraction of an average day 1kW of nameplate capacity can generate 1kWh of power, which for solar typically ranges on the order of 0.2-0.25, depending on geography. We will proceed in this paper by abstracting from the issue of

⁸Of course, it is worth noting that given enough time, the market would likely find a use for the excess power in the summertime, and it would not be wasted. Endogenously lower summer prices would send a signal to encourage flexible demand to ramp up in the summer. One can already imagine seasonal hydrogen fuel production, desalination, and flexible computation loads responding to lower summer electricity prices. We thus think of the 8-hour benchmark battery storage scenario we consider as fairly conservative. Tong et al. (2021) show that when optimally mixing renewable wind and solar energy production—even with a three-hour storage capacity—all the major economies in the world in terms of total GDP can offer grid reliability (share of demand met by supplied renewable electricity) of about 80% and upwards. Increasing the storage buffer to 8 hours

Figure 3: Solar Module Prices and Projections



Notes: Historical data is from Our World in Data, and is a composite of International Renewable Energy Agency (2023), Nemet (2009) and Farmer and Lafond (2016). “AW 2023” refers to projections in (Arkolakis and Walsh, 2023). “Trend: Full Sample” takes the average decline in the historical data and projects out prices from 2023 onwards. “Trend: 2000 on” does the same using the average decline in the post-2000 data. The datapoint for 2023 is not included in these projections.

seasonable variability. However, we found that applying the curtailment estimates from Arkolakis and Walsh (2023) to lower firmed solar’s capacity factor, as well as increasing the storage buffer from 8 to 12 hours, do not meaningfully change the results below.⁹

We now discuss how we project the current costs of firmed solar power out to 2040.

We begin with the cost of the panels. In Figure 3 we show the historical price of a solar panel per watt of output, which corresponds to the “Panel” cost in Figure 2. The decline in price for panels has been extremely fast and prolonged by any standard, averaging 11% annually in real terms since 1980, and 13% since 2000. This has caused the price of panels per watt to decline about 100-fold between 1980 and 2020. It is safe to say that no-one in previous decades imagined we would be here in 2024, with solar now the cheapest form of unsubsidized bulk energy supply in most parts of the world. The question that confronts us now is where this trend is heading.

There are a number of analyses in the literature that attempt to project solar costs into the future. One of these is our own (Arkolakis and Walsh, 2023), which uses a structural

⁹We find an increase of 20-30% in the wholesale price bounds across space in either of these exercises, which still implies significant price drops against current prices for most of the US.

model of the world economy’s adoption of renewables, where progress in capital costs is driven by “learning by doing”, and which is disciplined with grid-level parameters. Another influential piece in policy circles has been Way et al. (2022), which estimates statistical experience curves for a range of technologies, including solar, and backtests these against historical data. We plot both of these out projections out to 2040 in Figure 3. Both of these analyses predict a slowing in the trend rate of decline, and for largely the same reason: as solar’s share of the world generation mix expands to significant levels, the next doubling of capacity becomes progressively harder to achieve. While this is a reasonable assumption, both of these projections look to have already been proved too conservative by the stunning data in 2023 (visible on the graph), when costs declined almost 45% in a single year. Historically, one could have done much worse at any point in the last 40 years than just drawing a straight line in log space and pushing that forward a decade. Such “naive” forecasts are also included in Figure 3.

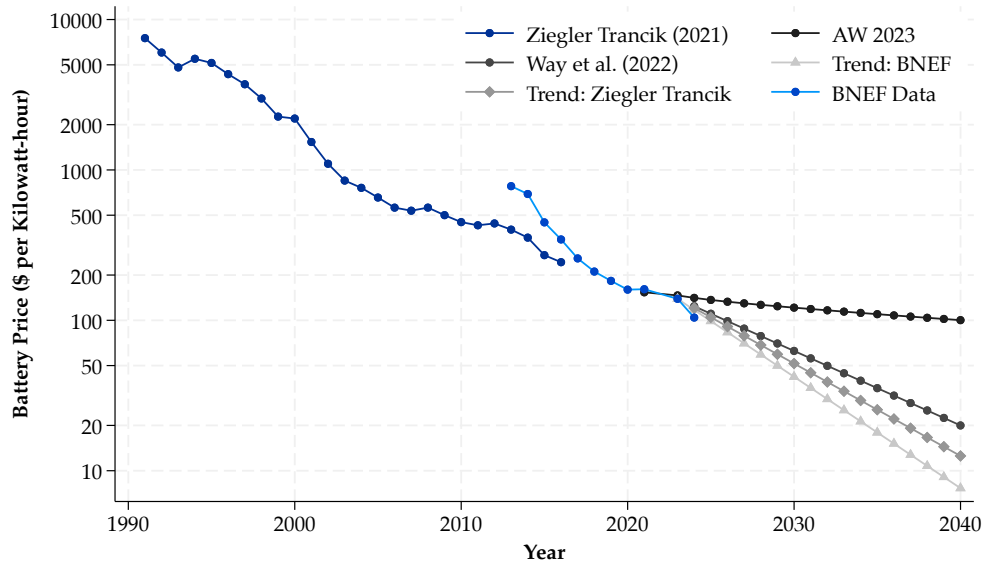
Projecting costs at the historical rate of decline leads one to incredibly cheap capital costs by 2040. Could a solar module really cost \$20 per kilowatt, or even \$10? This seems somewhat fanciful from our vantage point in 2024. The raw materials alone currently run to more than that, long before we think about the costs involved in production. At current prices, the silicon in a 700 W module costs around \$5. Then there is \$6 of aluminum for the frame, and \$9 dollars of copper for the wiring. This is to say nothing of the glass, which currently makes up 75% of the weight of a 35 kg 700W panel.¹⁰

But a little imagination can get us a long way from the \$145 of mid 2023 to \$20 in 2040. First, both research and commercial efficiencies of solar panels have been steadily improving for decades. Perhaps the most famous figure ever produced by the NREL charts the progress of record solar cell efficiencies by panel type (reproduced in Figure 20). Current commercial solar cell efficiencies are around 20%. Improving this to 30%, well within the realm of current lab efficiencies for multi-junction and hybrid cells, would lower cost by a third for the same materials. Stripping out raw materials, as would be possible in a move from monocrystalline to thin film or hybrid perovskite technologies, could lower costs by a similar magnitude. Due to fact that disparate improvements propagate multiplicatively to final cost, a further two-fold increase in manufacturing efficiencies, something that has been achieved many times in the last few decades, suddenly gets us in the ballpark.

For this exercise, we will take a middle road, and assume a cost decline that places the 2040 cost of solar cells halfway between naive trend extrapolations and the current vintage of model projections, so that a 2040 solar panel costs \$40 per kilowatt.

¹⁰Our estimates for these numbers use data from Dominish et al. (2019).

Figure 4: Battery Pack Prices and Projections



Notes: Data is from Ziegler and Trancik (2021) and Bloomberg New Energy Finance (BNEF). “Trend: Ziegler Trancik” takes the average decline in this data series and projects it out from 2023 to 2040. “Trend: BNEF” does the same using the average decline in the BNEF data. “AW 2023” refers to projections in (Arkolakis and Walsh, 2023).

Battery prices are today the largest component of a unit of firmed solar power. 8 hours of storage capacity from lithium-ion batteries currently adds \$1100 to the cost of a kW of solar power. While certainly expensive, this cost has been declining precipitously, by between 5 and 7 times in just 10 years, as shown in Figure 4. Going back 30 years, the cost of lithium-ion storage has fallen around 50-fold. This is around as fast as the price of solar modules have declined, and among the fastest cost declines recorded for any industrial good in the US. Assuming the trend continues for the next 16 years leads to around \$10 kWh. We will choose to be more conservative than these log-linear extrapolations, and use the numbers from Way et al. (2022), which leads us to around \$20 per kWh by 2040, i.e. \$160 for an 8-hour battery per kW. We note, however, that as with the huge decline in solar panel prices in 2023, battery prices have already diverged below these projections, falling 14% from 2022 to 2023, and a further 25% to mid 2024.

Balance of system (BOS) costs are the additional electrical components, such as transformers, module racks and inverters, that are needed to complete the installation and connect it to the grid. In historical forecasts of solar price declines, these costs were seen as a crucial bottleneck hampering continued price falls. In practice however, being mainly manufactured components, they have fallen quickly in price as well. NREL estimates that these declined by about 60% between 2010 and 2022, or a yearly rate of decline 8.4%. We will

assume this continues out to 2040.

Labor used in construction (“Fieldwork” in the terminology of NREL) consists of the labor required to mount, install and connect the panels to the grid. While falls in the cost of labor itself are unlikely, there has been consistent learning by doing in installation labor at both the utility and residential scale in solar that has improved overall construction efficiency in recent years. Then too, as the cost of panels and batteries decline, the economics of larger installations becomes more feasible, as some of this labor is a fixed cost and can be spread over more units.¹¹ These forces are difficult to size quantitatively. We assume that the cost of fieldwork per kW declines by a third by 2040, but little changes in our analysis if this component of cost remains constant.

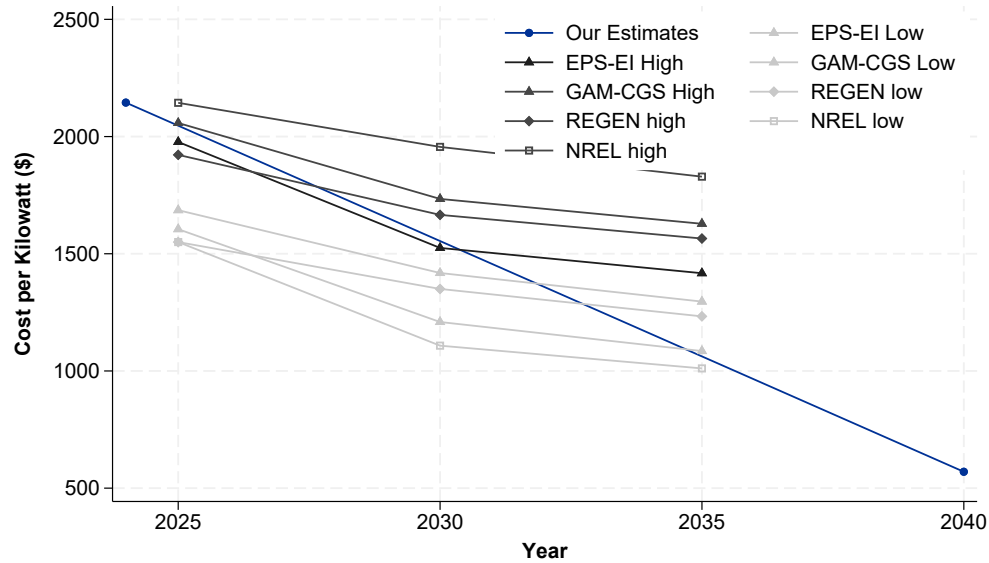
Land lease payments to host solar farms are currently a relatively small fraction of costs. This, however, may change in the future, as land prices are unlikely to fall significantly in coming years. As other components decline, land becomes increasingly important. Moreover, it is the only component of cost that differs significantly across space. We use data from Nolte (2020), who provides high resolution estimates of private land values. We average these at the county level (estimates presented in Figure 21). Land values differ significantly across the United States, ranging from \$1000 per hectare in remote rural areas to over \$1 million per hectare in New York City. These magnitudes significantly affect the viability of solar projects in densely populated areas. We assume that these land prices will be unchanged in real terms in 2040.

The “Regulatory” component computes the costs of applying for permits and environmental approvals, which we leave unchanged to 2040. NREL additionally lists a category of “Other” costs involved in project construction, consisting of Sales Tax, Management and Profit. Since conceptually these costs are percentage additions to the final installed project cost, we assume they remain proportional in 2040, and scale down accordingly. The final wholesale power price bound is presented in Figure 2. A unit of firmed solar power falls from \$2,145 in 2024 to \$570 in 2040.

Land, labor and regulatory costs eventually become the dominant component of installed capital costs, accounting for more than half of total cost by 2040. Beyond this point, the power of further falls in solar module, battery and balance of system costs to push down the steady-state wholesale price of power becomes muted. Indeed, without significant efficiencies in installation labor, and big increases in module efficiency that would allow

¹¹Automation of installation and replacing construction labor with robots is also a nascent possibility, see here.

Figure 5: Firmed Solar Cost Comparisons



Notes: The Figure shows the comparison of firm solar capital costs based on our calculations and different models in the literature, using data from the Appendix S1 of Bistline et al. (2023b). We extract solar PV estimates for the different models add low and high estimates of battery capital costs from estimates of NREL (Cole and Frazier (2023)) by multiplying their 4-hour battery costs estimates by 2, to obtain an 8-hour estimate. The models detailed here are EPS-EI, which refers to the Energy Policy Simulator, GAM-CGS the Global Change analysis model, the REGEN-EPRI the U.S. Regional Economy, Greenhouse Gas, Energy (REGEN) model developed and maintained by the Electric Power Research Institute (EPRI), and NREL the U.S. Department of Energy National Renewable Energy Laboratory’s Regional Energy Deployment System (ReEDS). Prices are in 2023 dollars.

a significant reduction in land costs, it is very difficult to see how solar project costs go below 200-300 dollars per kilowatt even in the very long run. We state this with some caution, of course, noting the long history of failed predictions and overages on how low solar costs could conceivably go.

Comparison of our estimates to leading projections in the energy systems literature discussed above reveals that our estimates are optimistic, but not unrealistic. In Figure 5 we juxtapose our predicted firm solar capital costs to estimates from various simulated energy system models in the literature, as summarized by Bistline et al. (2023b). We see that our estimates mostly agree with others in the literature in 2030, but then the decline in our estimates continues to new lows while declines in all other estimates tend to peter out from 2030 to 2035. At the heart of this difference is that all those models assume that both solar PV and battery cost declines will start significantly slowing down in the next decade. We also assume that the declines will continue at a slower pace than recent decades, but that this pace will still be quantitatively significant.

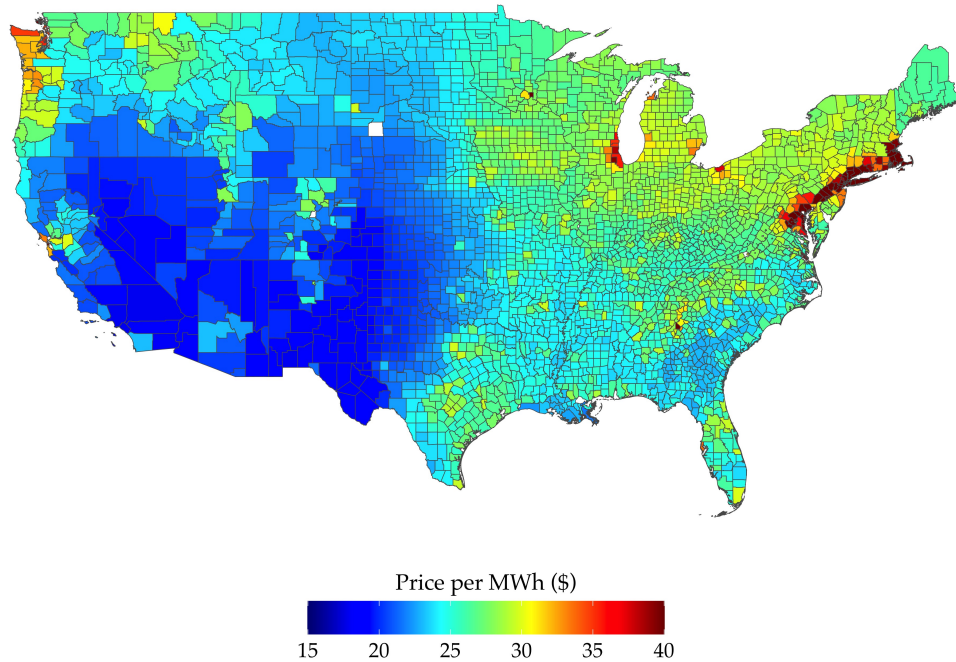
We now use our estimated 2040 capital cost in conjunction with equation (2) to compute the bound on wholesale prices in 2040. To do so, we first need estimates of θ_ℓ , the average amount of electricity produced by a panel in a year at different places in the US. This is provided by the Global Solar Atlas for the United States at a very granular level, and we plot estimates of panel output in Figure 18. The spatial variation is marked: a 1 kW system in Southern California produces 1,825 kWh a year. The same system in Seattle produces 1,277 kWh a year. This significantly impacts the future implied prices across space, since in sunnier areas it translates directly into a lower price per each kWh needed to justify the upfront capital cost of equation (3).

Lastly, we assume a long run financing cost of 5%, which is consistent with current estimates from the International Energy Agency for solar in the United States. Of course, this depends on a return to a relatively low interest rate environment in 2040; any further breakouts of inflation in the next decade would raise this number.

The final implied price bounds from equation (2) incorporating all this information are plotted at the county level in Figure 6. Prices range from around \$20 per MWh in sunny, sparsely-populated areas, to above \$35 in densely-populated urban corridors, with the payroll-weighted average being \$27. In most parts of the country, this represents a significant decline on wholesale prices in 2024. We collect current wholesale prices from the Energy Information Administration for the major Regional Transmission Organizations (RTOs) and Independent System Operator (ISO) pricing hubs, as well as regions that do not use locational marginal prices, and then plot the implied price decline in Figure 22. Price declines range from 20% in the densely-populated parts of Midwest, to 40% in the New York and the South, and all the way up to 80% in California, Texas and much of the West.

The advantage of taking this bounds approach is that we can say something about future prices with a minimal set of assumptions. As we will see below, knowing the bounds on electricity prices allows us to derive a full set of general equilibrium wage responses, and assess the macroeconomic impact of the renewable transition. The disadvantage is that we can say nothing about quantities. We cannot predict how much firmed solar capacity will be installed in any particular place, nor even a potential range of quantities. For that, one really does need a fully-specified structural model of the US economy, complete with assumptions about local demand curves for electricity in general equilibrium, and stocks of alternative technology capital like natural gas and nuclear (such as in Arkolakis and Walsh (2023)).

Figure 6: Future Implied Wholesale Price Bounds Across the US



Notes: Figure shows the implied bound on wholesale prices using equation 2 for 2040. Prices are in 2023 dollars.

Nevertheless, the advantage of our approach is significant, and so it is worth probing a little further the minimal assumptions the analysis above rests upon.

First, let us note that our methodology does *not* assume that the grid of the future is fully supplied by firmed solar. Even with large amounts of battery storage, there will still likely be a need for gas backup, in certain areas and at certain times, for years to come. Second, the fleet of US nuclear plants are fully depreciated at this point, and they have a minimal marginal cost of supply. While the existing plants are relatively old on average, there appears to be little technical barrier to extending their operating lifetimes for decades more.¹² The electricity they produce is carbon-free, and also free of the many additional pollutants pumped into the atmosphere by coal and natural gas turbines.¹³ As such, they

¹²See, for example, the Department of Energy here.

¹³Given that it is somewhat difficult to adjust nuclear power plant output rapidly, nuclear plants are particularly unsuitable to a regime dominated by unfirmed solar power, where high output during the day pushes down midday prices dramatically. In recent times California and Australia, with their excellent solar insolation, have seen protracted negative price events during the day, with generators having to pay to bid

are a valuable asset in a world of clean power, and will likely continue to supply up to 800 terawatt-hours (TWh) for many years to come, enough for 20% of current US electricity demand.

Instead, what we assume holds is a simple no-arbitrage condition: that the low fixed costs of firmed solar, and its relative ease of construction, places a ceiling on what local generators can charge in the medium term without inducing additional solar entry. Behind this is the implicit assumption that marginal entry into new solar within each county is elastic, so that if wholesale electricity prices rise (say, because demand increases), new firmed solar can easily be constructed locally. One might well ask if this is true in all areas in the US.

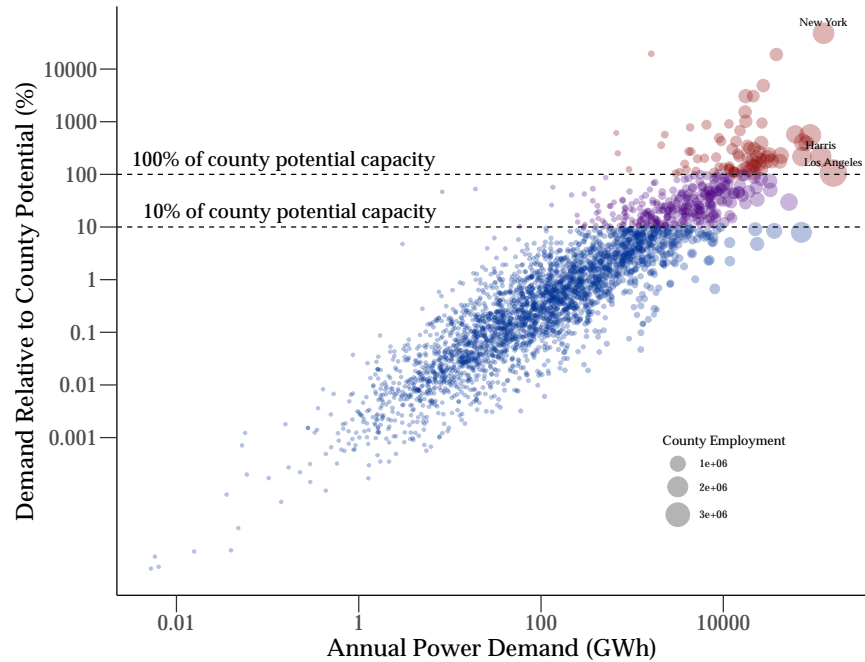
In particular, a dominant concern in popular analyses has been whether there is enough land for the renewable energy transition. Is there enough land available locally for elastic entry to be a reasonable assumption? In the aggregate, available land is clearly not a binding constraint. The US consumed 4,178 TWh of electricity in 2023 according to the Energy Information Administration (EIA). Using data from the National Renewable Energy Laboratory (NREL) on developable land area for solar power in the lower 48 states of the continental US, we estimate that there are over 83 terrawatts (TW) of potential solar capacity available in their “Reference Case”. Combining this with the data used above on total solar insolation at the county level, total current US power demand represents about 2% of developable solar capacity in the continental US.¹⁴

But what about locally? We can use data from the Quarterly Census of Employment and Wages (QCEW) to form a detailed picture of the county industrial structure at a relatively fine level of industry disaggregation (6-digit NAICS, or around 1000 different industries). Some industries, such as mining, milling and manufacturing, are relatively electricity intensive, using large amounts of power for production. Other industries, such as legal services and personal care, use comparatively little electricity. Differences in local industry structures then lead to differing amounts of power usage across space. We can get a measure of relative industry electricity demands from aggregate sectoral data in the input-output tables, which records how much each industry spends on electricity annually. Assuming this is approximately proportionate at the local level allows us to form a county-level estimate of demand for electricity from industrial and commercial use. To

into supply. To us, this seems like a clearly temporary phenomenon as battery storage scales up.

¹⁴Total developable solar capacity in the Reference Case of the NREL dataset represents around 30% of land area in the continental US. As such, meeting current demand solely from solar power is feasible using less than 0.6% of total land in the continental US, not inclusive of rooftop potential. For reference, this would be around 5 times the land currently used for golf courses.

Figure 7: Current County Demand Relative to Local Potential



Notes: Data on local potential is from the NREL estimates of developable capacity, in their Reference Case, combined with insolation data from the Global Solar Atlas. Local demand is estimated using the local county employment mix by industry from the Quarterly Census of Employment and Wages, combined with sectoral electricity usage in the 2017 detailed input-output tables from the Bureau of Economic Analysis.

complete the picture, we add residential demand, assuming this is proportional to county population.¹⁵

In Figure 7 we plot an estimate of county level demand for electricity in 2023 against the fraction of local developable solar potential capacity from the NREL data that would need to be developed to meet that demand. We use their Reference case to form county-level estimates of developable solar capacity. Under this case, many areas are excluded on a very fine geographic scale from developable potential under this case: built up urban areas, conservation easements, federal Department of Defense lands, infrastructure setbacks, regulatory bans and moratoriums, and elevated/unsuitable terrain.

We estimate that the vast majority of counties in the US would be able to meet their power needs locally, without the need for transmission from other counties. Of the 3,109 US counties, 2,705 could meet all of their current power demand by developing less than 10% of the land the NREL estimates could be converted to solar power production. We note that this is far less than 10% of the actual land of the county. Another 312 counties could

¹⁵While this will not be exactly true, as the local climate will impact electricity use per household, it is a reasonable first step.

conceivably run entirely on local supply, but would need to develop much or most of the potential capacity the NREL estimates is available.

However, 92 US counties use power far in excess of what could ever be supplied locally by firmed solar. These are mainly populous, urban counties such as Los Angeles County, Harris County in Chicago and New York County, which covers the island of Manhattan. NREL estimates that there are 3 square kilometers of Manhattan that could in theory be turned over to solar power production (out of a total of 59). This would provide power for about 0.02% of Manhattan's demand. Clearly, populous urban counties will need to source their power from other places.

But how much long-distance transmission is actually required? In Figure 8 we consider the amount of potential available in a broader local area: within 100 miles of the county centroid. We add up all the local developable solar potential, as well as all demand from counties within this radius, and then compute a measure of demand relative to potential within 100 miles. We plot this against the total land requirement of meeting demand. Now even the large urban counties (for the most part) are well able to meet their demand locally. Moreover, 94% of counties would require less than 5% of local land for power production, and in most cases significantly less.

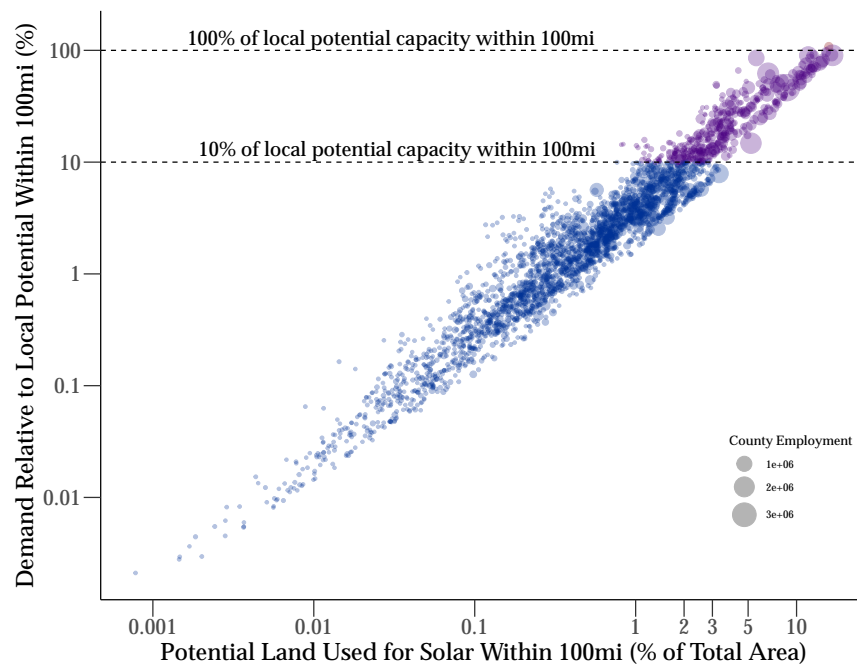
While we make no prediction that autarky will be the actual outcome for many counties, this analysis does support the assumption that elastic entry is not severely challenged by local constraints on the availability of developable land for solar projects. As such, solar capital costs will create competitive pressure on wholesale generation in all parts of the US. Meanwhile, the transmission network will remain crucial in ensuring access to low cost supply from other areas, particularly for dense urban areas. We return to considering the build out of the transmission network in Section 3.4.

3 The Macroeconomic Impacts of Lower Power Prices

Wholesale prices fall anywhere from 20% to 80% out to 2040, depending on the local solar resource, initial electricity costs and local land costs. This is a large change in a key input price into production. How should we think about the impact of such changes on economic output?

Much recent work in economics has studied the passthrough from fundamental productivity and price shocks through to final economic activity and welfare. Many considerations emerge, such as how the shocks affect market power, reallocation of factors across

Figure 8: Local Land Requirements for Solar



Notes: Data on local potential and land requirements is from the NREL estimates of developable capacity, in their Reference Case, combined with insolation data from the Global Solar Atlas. Local demand is estimated using the local county employment mix by industry from the Quarterly Census of Employment and Wages, combined with sectoral electricity usage in the 2017 detailed input-output tables from the Bureau of Economic Analysis.

uses, and the direction of technical change (see Baqaee and Rubbo (2023) for a discussion of some of these issues). Here, we will keep our analysis relatively simple, and focus on the regional exposure to energy price falls in general equilibrium. We conduct our analysis at the county level.

3.1 The Setup

Consider a spatial economy with locations ℓ , and sectors s . Workers live in these locations (for example, Los Angeles County CA, or Yellowstone County MT) and choose a sector to work in (for example, aluminum smelting, finance, or hospitality). Firms produce and sell products within narrowly defined sectors, selling their products both locally and across the country. To produce, they need to hire labor, buy some intermediates inputs, and use some electricity.

We assume that firms produce a unique differentiated variety i . Firm i located in location ℓ and sector s produces output according to the production function

$$y_i = z_i F_{\ell s}(l, e, X),$$

where l is labor, e is electricity and X is an aggregator of a vector of intermediate sectoral inputs \mathbf{x} . z_i is an index of firm level TFP as in Melitz (2003). Note that the production function $F_{\ell s}(\cdot)$ may be location- and sector-specific. This production function allows, for example, for exogenous local productivity differences as in Redding and Rossi-Hansberg (2017), and endogenous agglomeration forces that are location-sector specific, as in Bartelme et al. (2017), as long as these are taken as given by the firms.

Electricity is a local factor of production, with a price $p_\ell^\mathcal{E}$. Labor has a price specific to the location and the sector, $w_{\ell s}$, arising from imperfect substitution across sectors within a location, and less than infinite elasticity of labor supply across space.

The exogenous productivity z_i is drawn from a distribution $\Psi_{\ell s}(z)$, which may depend on location or sector. Intermediate inputs enter through the aggregator in a symmetric way for all firms (though they may use this with different intensity) so that $X = f(\mathbf{x})$. We further assume that this aggregator take the same form as the final good aggregator, so that both the intermediates price and the final good price serve as the numeraire.

Firms innovate a variety and enter the market by paying an entry cost, defined by $g_{\ell s}(l, e) = 1$, in terms of local labor and electricity. We assume g is constant returns to scale. The resulting entry cost is denoted $G_{\ell s}(w_{\ell s}, p_\ell^\mathcal{E})$. Firms exit at an exogenous rate ξ .

Cost minimization for a given level of output y allows us to write a cost function

$$C_{\ell s}(y; w_{\ell s}, p_{\ell}^{\mathcal{E}}, z) = z^{-1} y v_{\ell s}(w_{\ell s}, p_{\ell}^{\mathcal{E}}),$$

where v is the average unit cost function for a location-sector pair.

We suppose the output market gives rise to a concave revenue function for the firm that takes the form

$$B_{\ell s}(y) = D_s r(y), \quad (4)$$

where $r(y)$ is continuously differentiable and concave, and where D_s is an aggregate sectoral demand shifter. This shifter is understood to be a fully endogenous object in general equilibrium, and a function of all prices in the economy, but the firm takes it as given.¹⁶ As we discuss in the Appendix, many demand systems have a revenue function that takes this form, including the classic Constant Elasticity of Substitution (CES) demand system and single aggregator demand systems (see e.g. Arkolakis et al., 2012, Matsuyama and Ushchev, 2017).

So far we have not said much about the consumer side of the model, except the ones implied by equation (4), or how the investment costs to create firms are financed. In the Appendix, we discuss a general framework that will lead to this demand structure, as well as specifying how workers choose where to work and live, and a full dynamic structure for preferences. For now, all we require is that worker labor supply is increasing in the wage $w_{\ell s}$. Furthermore, different assumptions on the demand and market structure will result in a different form of the market shifter D_s . We make such assumptions explicit in Section 3.3 and discuss how they allow us to solve for the general equilibrium of the model.

Without presenting and defining the full equilibrium structure for brevity (we do so in full in the Appendix), the key equilibrium condition we use for analysis is the free entry condition, which allows for an intuitive treatment of the local incidence of electricity price shocks. This ensures that discounted expected profits equals the entry cost, and in a steady-state takes the form

$$G(w_{\ell s}, p_{\ell}^{\mathcal{E}}) = \sum_{t=0}^{\infty} R_{0 \rightarrow t} \left(\int_z \left[\max_y D_s r(y) - C_{\ell s}(y; w_{\ell s}, p_{\ell}^{\mathcal{E}}, z) \right] d\Psi_{\ell s}(z) \right). \quad (5)$$

¹⁶Notice that this is not a completely innocuous assumption. For example, in our context it implies the absence of trade costs or differential wedges across location-sector pairs.

where as above, $R_{0 \rightarrow t}$ is the cumulative interest rate.

3.2 Local Wage Responses to Changes in Electricity Prices

Let the local price of electricity be $p_\ell^\mathcal{E}$ in the long run. We use the bounds derived above in Section 2 as an exogenous change in the local price of electricity driven by uptake of firmed solar power across the grid.

Proposition 1 *Assume that the free entry condition (5) holds. Then the general equilibrium response of wages $w_{\ell s}$ in location ℓ and sector s in response to a local change to the price of electricity is given by*

$$d \log w_{\ell s} = -\frac{\Phi_{\ell s}^E}{\Phi_{\ell s}^L} d \log p_\ell^\mathcal{E} + \left(\frac{\Phi_{\ell s}^E + \Phi_{\ell s}^L + \Phi_{\ell s}^X}{\Phi_{\ell s}^L} \right) d \log D_s, \quad (6)$$

where $\Phi_{\ell s}^E$ is total local sectoral expenditure on electricity, $\Phi_{\ell s}^L$ is expenditure on labor, and $\Phi_{\ell s}^X$ is total expenditure on intermediate inputs. $d \log D_s$ is a measure of sectoral demand changes as electricity prices fall across the country.

The intuition for this formula is simple.¹⁷ Focus first on the term concerning $d \log p_\ell^\mathcal{E}$. If local costs of electricity fall, all else equal this causes local firms to become more profitable. In general equilibrium, this causes new firms to enter and incumbent firms to increase their labor demand and output, until that increase in profitability is eroded away by higher wages, and balance is restored. The strength of this effect is directly proportional to the intensity of electricity relative to labor in production. The intuition for the $d \log D_s$ term is similar. If aggregate sectoral demand increases (say, because of rising incomes), and firms are not competitive price takers, then firm level profitability will again rise. This necessitates an increase in the cost of local labor to balance out the increase in profitability, and more so in places where labor is a smaller share of local sectoral input expenditure.

It is also worth stressing how general this result is, and thus how suitable for analyzing both the aggregate and distributional consequences of the transition to clean energy. To size the first term on the right hand side of equation (6), we need to know nothing about elasticities of substitution between electricity and other inputs, either at the firm level or at the aggregate level. Indeed, so long as relationship (4) is satisfied we do not need to

¹⁷A similar formula is derived in Eckert et al. (2022) for declines in the investment price of capital, and the methodology derived there is the basis for our analysis here.

know the details of any of the production functions. Importantly, we need no knowledge of firm-level heterogeneity or the firm size distribution.

Likewise, we need know nothing about labor supply elasticities or the ease of reallocating factors across firms to derive this expression. All we need is that labor is not perfectly mobile (or infinitely elastic) across space in response to wage changes, so that there is an upward sloping labor supply curve for each region, and for each sector. However, the shape of this curve is unimportant. Lastly, there is no requirement that the economy be efficient or close to an efficient equilibrium.¹⁸

Sizing the effect of the second term on local wages requires making further parametric restrictions, and we return to this in Section 3.3 below. The magnitude of the effect demands on a parameterization of “aggregate demand externalities”, which are common in models of monopolistic competition, and the exact details are model dependent. It is also worth noting that this term can be exactly zero in a model of competitive, price-taking firms with constant returns to scale production functions.

The nature of the exercise we undertake is to use the price bounds developed in Section 2 to form an estimate of $d \log p_\ell^\mathcal{E}$, and then use equation (6) to trace through the general equilibrium impact on wages. We focus on wages first because of the excellent local data on sectoral employment and wages, available at very fine levels of geographic and sectoral disaggregation. Second, as long as the aggregate labor share is stable (or almost stable), and differences in goods prices across space are abstracted from, wage changes are a simple sufficient statistic for welfare.¹⁹

Direct Regional Effects. While it easy to derive the impact on local sectoral wages, estimating the impact on average wages in a location requires knowing how easy it is for labor to reallocate across sectors. Let $\mu_{\ell s}$ be the employment share in sector s in location ℓ .

¹⁸While this formula bears some superficial resemblance to Hulten’s theorem and related results, it is quite distinct, and arises from the basic requirement of zero expected profit after firm entry costs are paid. If firm production functions are constant returns to scale, and markets are competitive, $d \log D_s = 0$ and only the first term on the right hand side appears. A more general version with different factors is derived in Eckert et al. (2022), from which we take inspiration here.

¹⁹In the class of models we consider in Appendix (B), with CES sectoral demand, firms profits are proportional to sales. Thus the only source of instability in factor shares comes from the aggregate elasticity of substitution between labor and energy. However, given energy’s quantitatively small share, in practice this means the labor share of aggregate income is approximately stable.

We can write the change in average wages in location ℓ as

$$\begin{aligned} d \log w_\ell &= d \log \left(\sum_s \mu_{\ell s} w_{\ell s} \right) \\ &= \sum_s \frac{\mu_{\ell s} w_{\ell s}}{\sum_{s'} \mu_{\ell s'} w_{\ell s'}} (d \log w_{\ell s} + d \log \mu_{\ell s}). \end{aligned} \quad (7)$$

Now suppose that we look at relatively fine industry classifications, so that no one industry is especially large. Furthermore suppose that the long-run labor supply elasticity is the same across industries and constant at η (we provide a standard microfoundation in the Appendix). We then define a measure of exposure of local wages to the electricity price by combining (6) with (7) to obtain


$$d \log w_\ell = \underbrace{\sum_s \frac{\mu_{\ell s} w_{\ell s}}{\sum_{s'} \mu_{\ell s'} w_{\ell s'}} \frac{\Phi_{\ell s}^E}{\Phi_{\ell s}^L}}_{\Omega_\ell \equiv \text{Direct Exposure}} d \log p_\ell^\mathcal{E} + \underbrace{\eta \sum_s \frac{\mu_{\ell s} w_{\ell s}}{\sum_{s'} \mu_{\ell s'} w_{\ell s'}} \left(\frac{\Phi_{\ell s}^E}{\Phi_{\ell s}^L} - \sum_{s'} \mu_{\ell s'} \frac{\Phi_{\ell s}^E}{\Phi_{\ell s}^L} \right)}_{\text{Labor Reallocation Across Sectors}} d \log p_\ell^\mathcal{E} + \underbrace{\Gamma_\ell}_{\text{G.E. term}}. \quad (8)$$

We begin the quantitative analysis by examining wage growth induced by direct exposure. We calculate the direct exposure, Ω_ℓ , at the county level. To get measures of local payroll and employment, we use the Quarterly Census of Employment and Wages for 2023. We let s be a four-digit NAICS sector.

Electricity intensity is only measured well in a comprehensive way in the national input-output tables. For 392 four-digit industries, we construct the ratio of total sectoral expenditure on “Electric power generation, transmission, and distribution” to total employee compensation, which measures $\frac{\Phi_{\ell s}^E}{\Phi_{\ell s}^L}$ at the sectoral level. We then do two imputations. First, we impute this ratio for missing industries via averaging the ratio (weighted by employment) at the three digit level for the four-digit industries for which we have observations, and then applying this ratio to the missing four digit industries. We repeat the procedure at the two digit level for four digit industries which have no other observations in their three-digit family. Second, we use the national-level industry ratios to proxy for the local level ratios. We recognize that heterogeneity in local factor prices are likely to cause some measurement error here, but without better local data this is a good first step.

There is a large amount of heterogeneity in electricity intensity by detailed industry code. In Table 1 we show the 25 most exposed industries. Far and away the biggest consumer of electricity is aluminum smelting, where the ratio of electricity payments to labor compensation is almost one-for-one. This is due to the energy intensive nature of the electrolysis

Figure 9: Direct Exposure to Electricity Price Falls



Notes: Figure shows the calculated direct exposure measures Ω_l from equation (8) at the county level. Data uses the Quarterly Census of Employment and Wages (QCEW) and the Bureau of Economic Analysis' Input-Output tables for 2023.

process, which converts alumina into aluminum useful for production. However, many other manufacturing and resource extraction industries, such as cement manufacturing, pulp mills and metal mining, are also highly electricity-intensive.

We then plot the direct exposure index Ω_l in Figure 9 against area size as measured by employment. There is a clear negative correlation with population. This arises because electricity shares are lowest in service establishments, particularly in non-tradable services like Retail, Hospitality and Education. In the data, there is a well-known strong correlation between population density and the percentage of employment in services, and this shows up in the ratio of payments to electricity to payments to wages inferred from the input-output tables. Large cities (employment above 1 million) have an average exposure of 0.026. This doubles in counties with population under 10,000, due to their proportionally greater employment in manufacturing and resource extraction.

The labor reallocation term in equation (8) occurs when the average wage in an area shifts due to employment composition changing in response to industry-level wages moving. It is governed by the local labor supply elasticity η , which determines how labor moves between sectors in the short-run in response to within-location wage changes.²⁰ Estimates

²⁰It is still the case that across-location labor supply elasticities have no bearing on the first-order wage change formula developed in (8).

in the literature tend to place this number around 0.2-0.7 (Artuç et al., 2010). We will use the average elasticity for both college and non-college workers estimated in Eckert et al. (2022) (which one of us developed), and take a value of 0.5. We plot the resulting exposure terms with labor reallocation added against the original exposure terms in Figure 24. Doing so has the effect of slightly muting the direct exposure in places that have less employment in exposed industries, since average wages in these less exposed industries (like Business Services) tend to be higher on average than wages in manufacturing. The opposite occurs in the more exposed areas. Overall though, for reasonable values of the labor supply elasticity this does not change the picture.

As such, areas with lower population density are in theory most exposed to the coming impacts of clean power. The actual changes, of course, will depend on the interaction of exposure with the falls in average prices, which are spatially heterogeneous.


In Figure 10 we plot the county-level average wage changes implied by interacting the direct exposure and labor reallocation terms of equation (8) with the 2040 implied price falls from Figure 22. Projected wage changes differ markedly across space. As with exposure, there is a mild rural bias in wage growth. This stands in stark contrast to recent wage growth trends in the US, which have overwhelmingly been urban-biased since 1980 (see Eckert et al. (2022)). Overall, the impacts from direct exposure are relatively modest: the average payroll-weighted real wage increase is 1.6% for the US as a whole. However, there are several large cities and counties that see greater rises. Table 2 in the Appendix shows the top 10: Salt Lake City, Los Angeles and Dallas all see real wage increases of almost 4%. This owes mostly to the fact that given their excellent solar insolation resources, and high current power prices, they are projected to see substantial falls in wholesale power prices.

3.3 General Equilibrium Aggregate Demand Effects

In general equilibrium, there may be additional effects associated with aggregate demand expansion. In general, lower electricity costs will lead to greater output, which in turn raises demand for all the firms in the economy. To be consistent with free entry, wages will rise further than is implied by the direct impacts. As noted above, such effects do not appear in all types of models, which is why we began with the direct effects. In particular, if markets are fully competitive and production is constant returns to scale, as in a traditional analysis, then the effects estimated above are the true effects.

Let us continue to abstract from the role of intermediates and trade costs. We further assume that sectoral spending shares at the aggregate level are constant, and that firms

Figure 10: Wage Changes from Direct Exposure



Notes: Figure shows the calculated direct exposure measures Ω_l from equation (8). Data uses the Quarterly Census of Employment and Wages (QCEW) and the Bureau of Economic Analysis' Input-Output tables for 2023.

face CES demand. In that case, we show in the Appendix that the change in the sectoral demand shifter firms face is simply given by

$$d \log D_s = d \log Y - d \log P_s^{1-\sigma_s}, \quad (9)$$

where Y is aggregate income, and P_s is a sectoral specific price index, with σ_s being the elasticity of substitution across goods. The intuition is straightforward. When electricity prices fall, aggregate income rises (because of greater output). All else equal, this increases spending on firms' goods coming through $d \log Y$. Offsetting this is that greater aggregate output induces entrepreneurship and the entry of new firms, who create more competition for existing firms. This shows up in the sectoral price index, $d \log P_s^{1-\sigma_s}$, and acts to dampen firm-level demand. Deriving expressions for these two objects is relatively

straightforward under standard assumptions on production functions.²¹ We have

$$d \log Y = \sum_{\ell} \sum_s \frac{\sigma_s}{\sigma_s - 1} \frac{(w_{\ell s} L_{\ell s} + p_{\ell}^{\mathcal{E}} E_{\ell s})}{Y} \left(\frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} (d \log w_{\ell s} + d \log L_{\ell s}) + \frac{\Phi_{\ell s}^E}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} (d \log p_{\ell}^{\mathcal{E}} + d \log E_{\ell s}) \right). \quad (10)$$

So the percentage change in aggregate income is just the activity-weighted change in local sectoral payroll and electricity sales.²² It turns out that under the same assumptions

$$d \log P_s^{1-\sigma_s} = d \log Y_s - \frac{1}{\sigma_s} \sum_{\ell} \frac{w_{\ell s} L_{\ell s} + p_{\ell}^{\mathcal{E}} E_{\ell s}}{\gamma_s Y} \left(\frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} d \log w_{\ell s} + \frac{\Phi_{\ell s}^E}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} d \log p_{\ell}^{\mathcal{E}} \right),$$

where the change in sectoral income $d \log Y_s$ is analogous to equation (10). Then, as long as we have some notion of the elasticity of substitution across firms at the sectoral level σ_s , we can compute these terms given the data we've already outlined above.

Directly estimating sectoral elasticities of substitution is complicated, and beyond the scope of this exercise. Many studies in the literature find values in the range of 3 to 8 (see, for example, Hottman et al. (2016); Gervais and Jensen (2019)).²³ We will use a value of 4 across industries, which is common in models of consumer demand and firm dynamics (Garcia-Macia et al., 2019; Peters and Walsh, 2022). One can show that a higher value for σ dampens the general equilibrium effects, and in the limit as $\sigma \rightarrow \infty$ there is no aggregate effect of demand expansion. In addition, we'll assume that the medium-run elasticity of demand for energy is around -0.5, consistent with estimates from the empirical literature (Labandeira et al., 2017).

In Figure (11) we show the county-level wage changes now including the general equilibrium effects. In general, these effects operate to dampen the wage changes of the most exposed places, as such places see more entry and firm creation (operating through $d \log P^{1-\sigma_s}$, which acts as a competitive spur to incumbent firms. As such, their profits increase by less than that implied by just the direct effect, and the local sectoral wage need not rise as much.

²¹In particular, that production functions are constant returns to scale, and that the costs to start a firm are denominated in units of the final good.

²²This arises because of the fact that in models of CES demand with constant returns to scale production functions, profits are just proportional to sales, and so are also a constant fraction of expenditure on inputs.

²³Demand elasticities, which correspond to the elasticities of substitution here, can also be inferred from estimates of markups at the industry level, such as in Hall (2018), which would give numbers around 4.

Figure 11: Wage Changes with General Equilibrium Effects



Notes: Figure compares the county level wage changes from direct exposure with the wage changes including the general equilibrium effects compute from equation (9).

In contrast, places with low direct exposure see higher wage growth. This mainly comes from aggregate income rising as power prices fall, some of which then gets spent on low-exposure industries like personal services, food and accommodation. One can think of the general equilibrium effects as redistributing the income gains from the most exposed places and sectors to the least. All in all, these effects are relatively modest, and do not substantially change the conclusions of the initial analysis. The aggregate effect on national wages from transitioning to firmed solar power rises from 1.8% to 2.6%.

The heterogeneity in the local wage responses also matter for general equilibrium aggregation. For the change in aggregate income in equation (10), the positive covariance between local expenditure shares and local electricity price drops amplifies the general equilibrium rise in aggregate income. This would be missed by deriving an aggregate version of the free entry equation without spatial and industry variation, and points to the importance of a disaggregated analysis for this question.

3.4 The Gains From Grid Integration

As we have emphasized, the spatial heterogeneity in price bounds presented in Figure 6 will not represent the true heterogeneity in prices observed on a solar-dominated grid. While the free-entry condition implies an upper bound on prices, actual prices will be sig-

nificantly below this bound in many areas. In particular, dense cities and suburbs, lacking cheap land for solar installations, will import much of their power consumption from surrounding areas using existing transmission infrastructure. The prices observed there will be closer to that seen in price nodes in rural areas, with adjustment for congestion.

The pricing formula used in many areas that implement Locational Marginal Pricing²⁴ is

$$p_\ell^\mathcal{E} = \underbrace{\mu}_{\text{System Generation Cost}} + \underbrace{\mu \frac{\partial L}{\partial D_\ell}}_{\text{Loss Adjustment}} + \underbrace{\sum_k z_k}_{\text{Transmission Constraints}} .$$

That is, the price in an area depends on three components. First is the system generation cost, or the marginal cost of generation for the last unit that bids into supply within that area. Second is a term that adjusts this cost by the marginal impact on system units of an additional unit of demand D_ℓ for power in location ℓ . Lastly, for each transmission line k that connects to location ℓ , there is an addition to the price which reflects whether that line is constrained, and the effect of an additional unit of demand at ℓ on the load on line k .

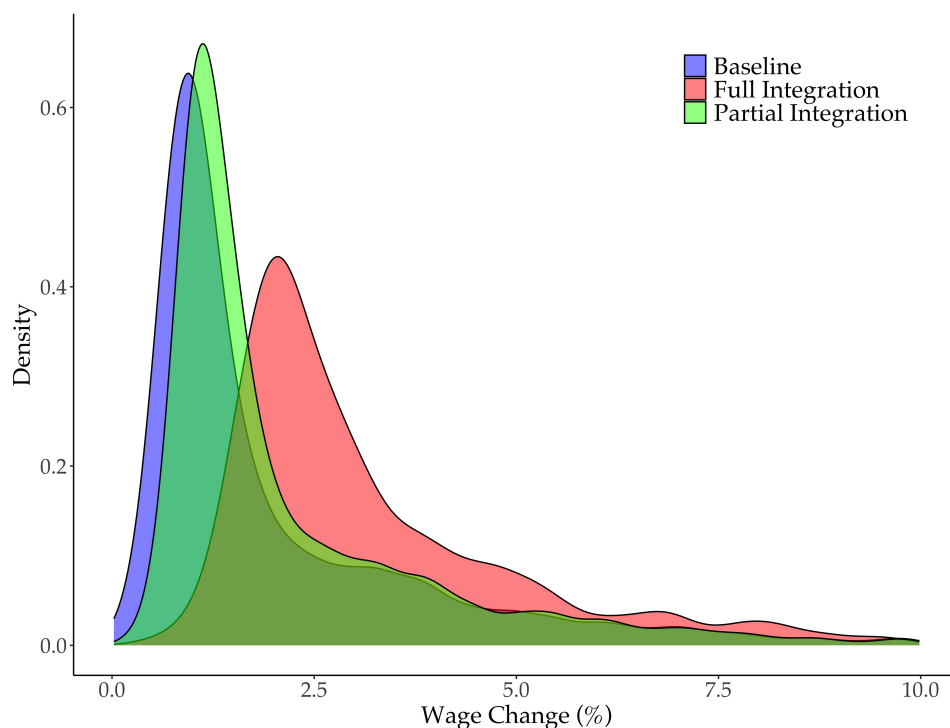
In a solar-dominated grid, the system generation cost μ will correspond to the average cost of generation of a unit of firmed solar power in the marginal areas connected to that Independent System Operator (ISO), as presented in equation (2).²⁵ For dense areas like New York County, which are unlikely to host solar farms within the city limits, the price of power will be determined by the generation cost in the rest of the (lower land cost) NY-ISO, along with adjustments for congestion and transmission. As such, the true variation in prices in the future is likely to be lower than implied by Figure 6.

But how much lower, and how costly is the resulting spatial heterogeneity in prices? Analyses of building out extra transmission capacity consider the effect of alleviating the congestion terms z_k , allowing high cost areas to take advantage of lower and lower cost marginal generating units in other areas (and potentially other ISOs when considering inter-organization flows). Doing so is a complex endeavor, as it requires solving a high-dimensional non-linear optimization problem. Here we come at the problem from a different angle.

²⁴See, for example, the NY-ISO.

²⁵We deliberately say “average” here, instead of marginal as the theory would imply. The marginal cost of generation of solar is effectively zero, and in an entirely solar-dominated grid the system generating cost must be the average cost of generation inclusive of capital cost for the most expensive solar unit, otherwise in the long-run capacity would exit.

Figure 12: Gains From Grid Integration



Notes: This Figure shows estimates of the gains from grid integration in two scenarios: “Partial Integration”, where enough capacity is built for the price to fall to at least \$25 MWh at all points in space, and “Full Integration”, where price becomes uniform across the grid at the lowest price in the US (around \$17 per MWh). Baseline refers to the wage changes computed using only the price bounds from Figure (6).

The theoretical maximum increase in production in the medium term from transitioning to a clean grid would occur if all locations could access Arizona’s generation cost of \$17 per MWh. Such an integrated continental grid is likely to be technically infeasible, even with huge investment in interstate transmission. However, it does serve as a way to size the prize on offer.

In Figure 12 we show the distribution across counties of the gains estimated using equation (8). In red, we recompute these gains if the implied price fall takes all location to \$17 per MWh. The economy wide payroll-weighted average wage increase rises from 1.6% to 2.7%. In addition, almost 20% of counties now see a wage increase in excess of 5%, particularly counties in the industrial Midwest and New England. To put this in perspective, assuming a stable labor share, this implies an additional increase in GDP of \$245 billion annually. To put it mildly, this is a large gain from a potential one-off investment.

However, achieving this gain really does require that all locations have access to power costs that are only achievable in the West of the continent. If we consider a more mild in-

tegration, the gains are much more muted. Suppose that sufficient integration is achieved to take the maximum wholesale power price across the US to \$25 per MWh (around the median price in the analysis above, and near the minimum in the eastern half of the continent), and less if the implied local power prices are below this threshold. This is represented by the green distribution in Figure 12. In this situation, the gains above the baseline are much more muted, rising from 1.6% to 1.9% in the aggregate. This suggests that the aggregate gains from harmonizing power prices in a renewable-dominated world may not be that large once solar determines the local price of generation. Further work on this topic is necessary.²⁶

How costly is it for the Northeast to gain access to the lower wholesale costs in the center and West of the continent? This is a difficult question to answer at a system-wide level. Conceptually, there are two issues to consider. First, how costly is it for firmed solar projects to connect to the high-voltage transmission network locally, so that the electricity they produce can be sent long-distance across the country? Second, how much does the capacity of the system as a whole need to be upgraded to handle greater cross-regional flows?

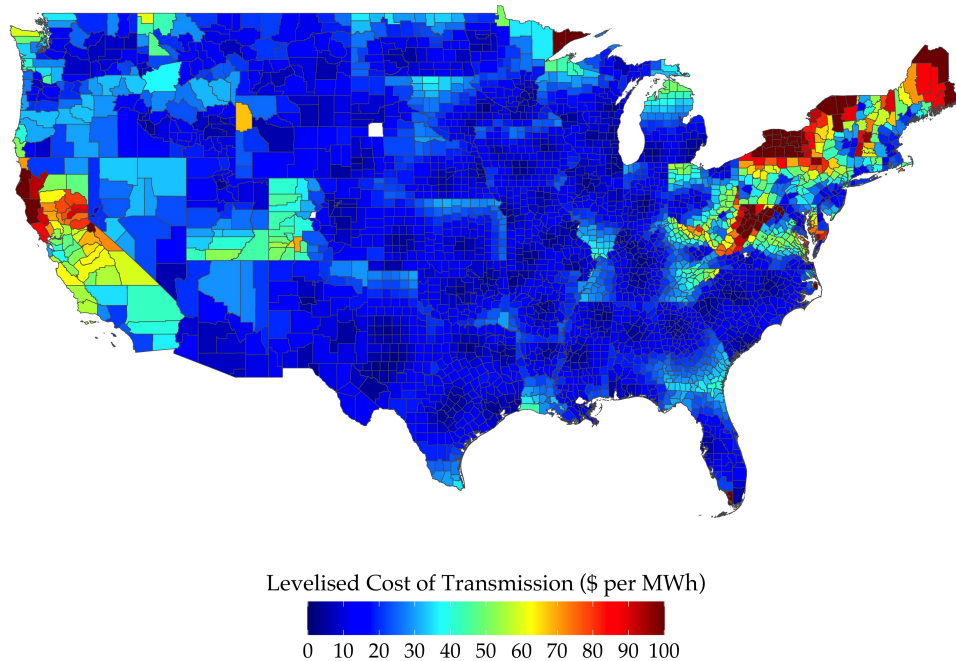
While the second would seem to require a fully-specified model of the US grid, for the first we can again use data from NREL, who construct localized estimates of the cost of connecting to transmission networks at fine geographic scale. This includes both the cost of building a “spur” connection from the solar site to the nearest substation on the electricity grid, as well as upgrading the substation and the local transmission network to handle the extra load.²⁷ In Figure 13 we plot this data at the county level.

Several patterns are apparent. First, transmission and interconnection costs for solar are low in much of the country. In particular, the low generating costs areas of West Texas, New Mexico, Colorado and Kansas also have very low costs of transmission, on the order of a few dollars per MWh. Particularly around established populations centers, these costs become negligible. Comparing with our results for the wholesale price bounds in Figure 6, there are large swathes of the Southwest that can feed into the grid with both minimal transmission connection costs, and low future wholesale prices. In contrast, connecting

²⁶See Gonzales et al. (2023) for an important real world case study of the effects of building out transmission lines in Chile to connect high insolation areas with dense urban loads. The authors find significant reductions on electricity prices in the wholesale market as a result of grid integration and entry of renewables.

²⁷The first element is also counted in the BOS component of the analysis of Section 2, and as such we do not try and add the total levelized cost of transmission to our estimates of the capital cost in Section 2. The second component is difficult to conceptually allocate entirely to the solar project developer, since network upgrades additionally benefit all other local projects and end consumers.

Figure 13: US Transmission Costs



Notes: This Figure shows average levelized cost of transmission at the county level. Data is from the National Renewable Energy Laboratory.

the quality solar resources of California to the grid appears to be quite costly. According to NREL, these differences are driven by regional construction costs multipliers for California and the Northeast, which drive up the cost of construction relative to the rest of the country. Permitting and environmental approvals are particular levers that could be examined to reduce these costs.

An additional concern is the effect of time delays in the interconnection queue, which is a common complaint of renewable project developers in 2024. The Lawrence Berkeley National Laboratory estimates that there are 2.6 TW of new capacity proposals waiting in the queue to receive approval to connect to the grid. Over 75% of these requests are solar, battery storage, or hybrid plants with both solar and storage. Depending on the region, getting approval often involves simulation studies of the effect of the project on local power flows and reliability. A key issue with renewable interconnections is that because of their smaller average sizes than the fossil fuel projects of the past, connecting renewable projects to the grid in tandem requires a greater number of reliability studies,

which appears to be significantly slowing approvals. Projects are now waiting up to five years to receive approval, up from two years in 2008.

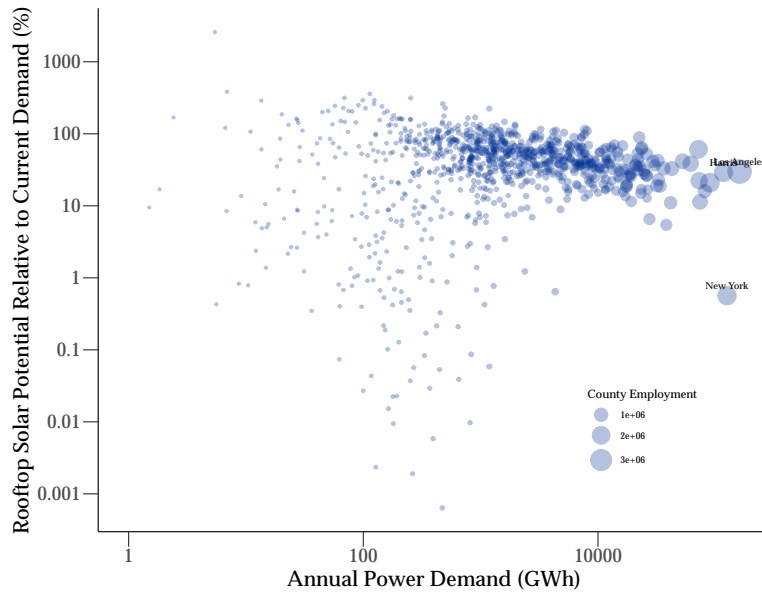
At one level, in our framework this is a transitional issue. Shifting the expected revenues even 5 years ahead into the future in equation (2) has a quantitatively small impact on the wholesale price bound for 2040, given relatively low prevailing interest rates. Through this lens, it has little impact on the macroeconomic impacts of clean power. At another level, however, increasing the expected wait time before interconnection can significantly increase risk for a project: risk of financing challenges, risk of regulatory changes and risk of development objections, to name a few. Project risk is not adequately captured by our framework. Nonetheless, it seems clear that such long lead times for interconnection are not necessary, and are a result of processes that could be streamlined. The Electric Reliability Council of Texas (ERCOT), for example, processes interconnection quests in under 2 years, more than 50% faster than the current mean wait of 33 months outside.

3.5 Passthrough

When we consider the effect of declines in power prices on real wage growth, we are assuming a one-to-one long run passthrough from wholesale to retail prices, which is unlikely to be the case in reality. According to the EIA, average US retail prices break down into 62% generation costs, 12% transmission costs (the cost of moving electricity down high-voltage transmission corridors) and 26% distribution (delivery to the final end user). While large industrial users typically pay close to wholesale generation prices, since their steady load and higher-voltage requirements limit the need for distribution infrastructure, residential users and smaller commercial users typically pay higher costs that include distribution charges as a markup on the wholesale cost of supply. Our arguments above mainly concern lowering generation costs; technical scope for reducing transmission and distribution costs seem more limited. However, it is worth making two brief points before concluding.

Much of the pricing model of regulated distribution networks reflects monopoly rents to “cover” the cost of distribution assets which have been fully depreciated. Many regulated utilities in the US (for example, PG&E operating under the California Public Utility Commission) operate via making a case to a government regulator for a “revenue requirement”, to cover the cost of owning and maintaining the distribution network, along with a regulated rate of return on their fixed asset base. This determines the markup on average wholesale power prices that they charge to the end user. It almost goes without saying that this is not a pricing model that has historically encouraged efficient investment.

Figure 14: Rooftop Solar Potential



Notes: Figure shows an estimate of rooftop solar potential at the county level from Google's Project Sunroof against our estimate of county demand. County estimates are missing in the Project Sunroof data for 2103 counties.

There has historically been no way for commercial and residential consumers to avoid joining these networks, as they act as natural monopolies. The arrival of modular rooftop solar with cheap storage changes this picture. Residential and commercial adoption of firmed solar power is likely to act as a competitive spur to regulated distribution monopolies, in a way that may encourage lower residential markups. While we may never see mass grid defection, with large numbers of users disconnecting from the grid altogether, the arrival of competition at the end use is a separate source of pressure on electricity prices, distinct from the lowering of wholesale generation costs.

Indeed, there is a large enough amount of rooftop solar potential in many counties to make this competitive possibility a threat to the distribution monopolies. In Figure 14 we use data from Google's Project Sunroof to form an estimate of the potential production of electricity on a county's roofs, and show this relative to our estimate of current county demand. In almost all counties for which we have data, rooftop potential ranges between 10 and 100% of current county demand. A notable exception is Manhattan, with its unparalleled density of skyscrapers having little relative roof space.

So far, rooftop solar uptake has been slow in the United States, given the dramatic cost falls in the price of panels that has driven utility-scale solar. NREL puts the cost of deploying residential solar in 2023 at more than twice the cost per kilowatt of utility scale

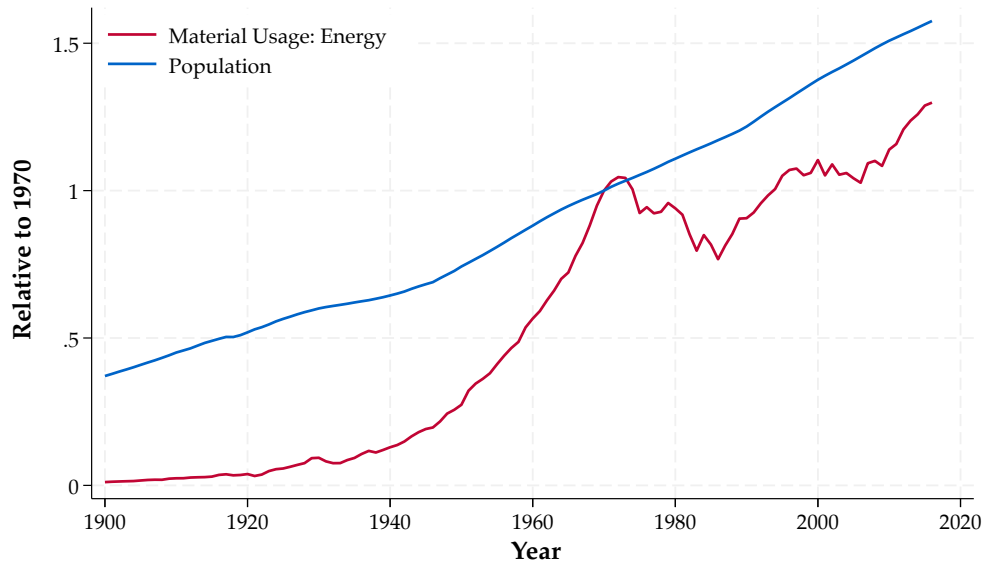
solar. Partly this is due to higher labor installation costs when the scale economies of utility installation are absent. However, the greatest difference in cost come from the soft costs of permitting and getting approval for the panels, which made up 29% of the total cost in 2023. This cost wedge is by no means an immutable fact of life: Australia, with insolation and incomes similar to California and Texas, now has solar panels on one third of the country's roofs, generating up to 10% of the country's electricity supply. Notably, installation costs are much lower than the United States, while the modules and equipment in both countries are common. This suggests that state- and city-level permitting reform to streamline installation processes could yield large benefits, not only through directly lowering solar costs, but by placing competitive pressure on distribution monopolies for the first time since the early days of the electricity grid.

3.6 Electrification

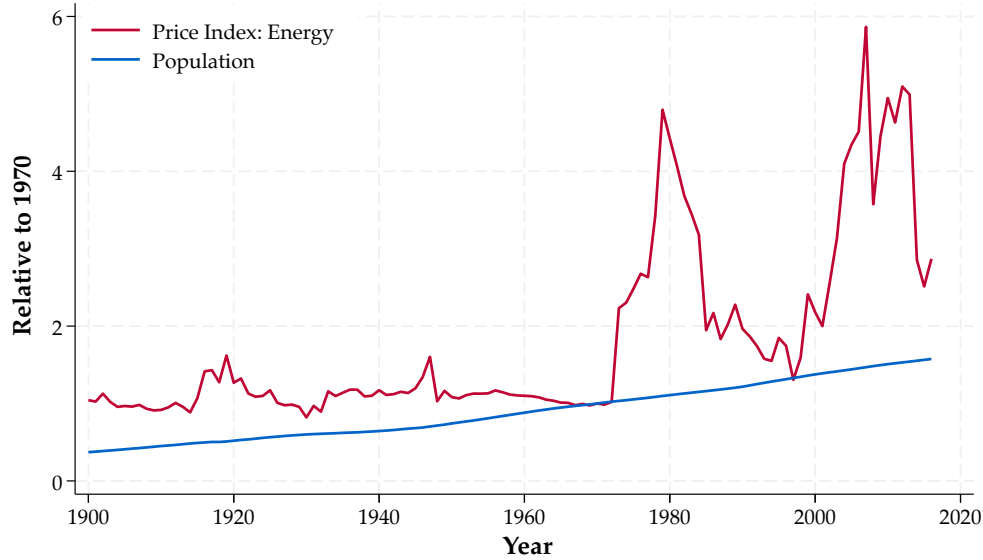
Our analysis makes no prediction about the quantities of electricity supplied by renewables; it is entirely price-based, which is both a strength and a weakness. As such, it is ill-suited to consider an increasing share of electricity in final production, which would have implications for the welfare impact of moving to clean energy. A big part of decarbonization efforts will involve electrifying significant parts of the economy. This includes shifting the transportation sector to primarily electric vehicles (a shift that is well underway worldwide), and moving building heat from oil and gas to electricity. The NREL projects that electrification could increase US electricity demand, which has been flat for decades, by between 20 and 38%. In addition, new sources of demand may further increase demand for electricity and its share in final production. The Department of Energy, for example, projects that increasing demand for electricity from data centers and AI could raise demand by an additional 20%.

A 40-60% increase in electricity output may not be associated with a significant increase in the electricity share $\Phi_{\ell s}^E$, especially as this analysis predicts that electricity prices fall by between 20-80%. The exact details will be location and industry specific in a way that is difficult to know before the fact. However, to the extent that the electricity share of final value-added is set to increase in a clean future, the welfare analysis in this section understates the true real wage gains from moving to a clean energy future.

Figure 15: Energy Inputs and Prices



(a) Inputs



(b) Prices

Notes: Panel (a) shows an index of total petroleum, coal and natural gas usage in the US economy against US population. Both index and population are shown relative to 1970. Data from 1949 onwards is from the Energy Information Administration, and pre-1949 is from digitizing Potter and Christy Jr (1962). Panel (b) constructs a Tornqvist price index for these commodities, divides this index by the CPI for urban consumers (CPI-U), and normalizes the resulting series relative to 1970.

4 Beyond The Medium Term: Removing Energy's Drag on Growth?

So the medium term impacts of the transition to clean power are likely to be relatively modest wage increases, with a moderate rural bias. Going beyond first order could take us into considering reallocation of factors across industries and space, to more energy intensive sectors, and to cheaper power locations.

However, there is another dimension in which the replacement of fossil fuel energy with clean technology based on manufactured goods represents a qualitative change in the structure of the economy. What the replacement of fossil fuels with clean electricity really does in the long run is turn energy from a problem of *finite resource extraction* to *capital accumulation*. This has significant implications for the direction of innovation at the aggregate level.

A number of analyses have pointed out that a structural shift occurred in the United States around 1970 related to energy use. Energy inputs from three primary sources, petroleum, coal and natural gas, had been growing strongly for decades. At the same time, prices for these inputs remained relatively stable in real terms. Then with the accompanying Arab oil embargo, the price of these inputs shot up, as can be seen in Figure 15. In subsequent years, these prices have stayed high, and have even been increasing on average relative to the pre-1970 period (though with considerable volatility). At the same time, energy usage has grown slower than population, a dramatic reversal from the pre-1970 pattern. As a result, energy intensity (in terms of joules per \$ of GDP) has been falling for many decades.

Hassler et al. (2021) show that this has been achieved through significant directed technical change in energy usage, with the productivity of energy usage in particular tripling since 1970 (after being roughly constant beforehand). It is easy to think of examples of greater energy efficiency in every day life. Fuel economy for vehicles has risen dramatically, lighting with LEDs uses an order of magnitude less electricity than incandescent bulbs, and household appliances like refrigerators and air conditioning consume much less power than their counterparts from 1980 and 1990.

These efficiency gains did not fall like manna from heaven; they were achieved through the purposeful use of innovative inputs like scientific labor and firm R&D spending. Given at any point in time, resources available for innovation are limited in the aggregate, this implies a tradeoff between energy-specific innovation and other forms of innovation.

Indeed, this shift in the direction of aggregate innovation may partly explain the relatively slow labor productivity growth observed since 1970. Nordhaus (2004) traces the sectoral propagation of the productivity slowdown in the 70s, and shows that it was concentrated in the most energy intensive sectors (see also Nordhaus et al. (1992) for a broader discussion of the “resource drag” in the pages of this journal).

The need to combat rising energy prices through greater energy efficiency investment acts as a drag on aggregate income growth. This drag can be completely eliminated in a world where energy production arises from capital accumulation, instead of extraction of a scarce input. To see why, we compare the Hassler et al. (2021) model of endogenous growth in an environment of resource scarcity with the same model in an environment where energy can be produced by accumulating renewable capital.

In the Hassler et al. (2021) model, there is a fixed amount of innovative resources that can be directed at improving productivity of either a capital and labor bundle, or energy inputs, both of which are needed to produce output. In most other respects, it is identical to the optimal growth model: capital can be accumulated, and the representative agent is maximizing intertemporal utility but choosing savings, consumption, and the allocation of innovative resources.

We present the model in the left column below.

Growth With Resources

$$\max_{c_t, k_{t+1}, e_t, n_t \in [0,1]} \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma},$$

subject to

$$c_t + k_{t+1} = F(A_t k_t^\alpha, A_{et} e_t) + (1 - \delta)k_t,$$

and

$$\frac{A_{t+1}}{A_t} = f(n_t) \quad \frac{A_{et+1}}{A_{et}} = f_e(1 - n_t),$$

and where resources e_t are finite, so

$$\sum_{t=0}^{\infty} e_t = E_0.$$

Growth with Renewables

$$\max_{c_t, k_{t+1}, k_{t+1}^R, n_t \in [0,1]} \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma},$$

subject to

$$c_t + k_{t+1} + k_{Rt+1} = F(A_t k_t^\alpha, A_{et}(\theta k_{Rt})) + (1 - \delta)k_t + (1 - \delta_R)k_{Rt},$$

and

$$\frac{A_{t+1}}{A_t} = f(n_t) \quad \frac{A_{et+1}}{A_{et}} = f_e(1 - n_t).$$

The notation is standard, but briefly, c_t is consumption, k_t is capital per worker, e_t is energy inputs (meant to represent exhaustible fossil fuels), A_t is factor-augmenting technical change and A_{et} is energy-augmenting technical change. n_t is innovative resources that can be directed either at improving factor-augmenting productivity A_t , or energy productivity A_{et} .

Importantly, energy inputs are in finite supply. Without innovation in how productive these inputs are, and similarly no innovation in capital and labor productivity, this is akin to the classic cake-eating problem that is used to teach dynamic optimization, and our diminishing supply of finite resources causes consumption to diminish over time. This is also one way to think about a model of “degrowth”, as advocated by Hickel (2021) and others.

Even with continued growth in factor-augmenting productivity A_t , it is not a given that growth in consumption is possible in the long-run without growth in A_{et} when resources are in finite supply. As pointed out by Solow (1974) and Stiglitz (1980), if the elasticity of substitution is greater than or equal to 1, either in a CES formulation or asymptotically for $F(\cdot)$ as $e \rightarrow 0$, then growth in the long run is possible. Greater technical progress and

accumulation of capital can offset the diminishing supply of energy. If, however, energy and production factors are globally complements, then the long-run path for the global economy features falling resources use, falling consumption and a rising resource share.

In the formulation of Hassler et al. (2021), innovation can be directed towards improving energy efficiency, and this alleviates this restriction. No matter the elasticity of substitution, we escape the curse of finite resources. The balanced growth path (BGP) solution involves resource use falling at the constant rate $\beta g^{1-\sigma}$, where g is the growth rate of output, and an ever-rising shadow price for the resource as it is used up.²⁸ Innovation is positive in both factor-augmenting productivity A_t , and energy productivity A_{et} , and innovative resources are directed to both sectors. However, resources are still a “drag” on aggregate growth.

In the right-hand column, we modify the economy so that energy is instead produced by capital. Recall that the first two fundamental features of renewables we discussed at the start of this essay were that there were modular (so that power output is a linear function of capital installed), and that they had zero resource cost. As such, we substitute exhaustible fossil inputs e_t for renewable capital k_{Rt} , which produces energy at rate θ (as in the analysis of the preceding section). We’ll also assume for simplicity that there are no exhaustible resources used in the production of renewable capital, silicon and iron being in such abundance in the earth’s crust that they are not worth modeling.

One can show that the balanced growth path in this economy looks quite different. With this small modification, there is no long-run improvement in energy efficiency, and all innovative resources are deployed to factor-augmenting technical change. In a sense, renewables remove the scarcity of fossil fuels, as energy produced is limited by the amount of capital that can be accumulated, not by how much of a finite resource can be extracted.²⁹ The long-run growth in output g is also strictly higher than that in the Hassler et al. (2021) world of the left hand side.

This result is reminiscent of the celebrated Uzawa (1961) theorem, in which all technical progress in the long-run must be labor-augmenting (see Jones and Scrimgeour (2005) for

²⁸Decentralizing this economy can be done in a straightforward way by modeling the incentives of private firms to invest in the two kinds of technical change. Of course, there is no guarantee that the competitive economy is efficient.

²⁹Given our discussion of land requirements for solar above, it is worth noting that it is not strictly speaking true that there is no resource requirement for renewables. The quantity of land is, after all, finite, and places a ceiling on the amount of renewables that may be installed. However, given how much solar potential there is relative to current US power demand (see above), we believe this ceiling is comfortably far enough for this to be a meaningful abstraction.

a discussion). The basic intuition of the Uzawa theorem is that because capital is accumulated, and labor is not, the trend in capital inherits the trend in total output. “Effective inputs” have to grow at the same rate for factor shares to be stable, so effective capital (capital multiplied by capital-augmenting productivity) has to grow at the rate of effective labor (labor multiplied by labor-augmenting productivity). Because capital alone is growing at the rate of output, “effective capital” must also be, and this has to be equal to the growth rate of labor-augmenting productivity.


Something similar happens here. For long-run factor shares to be stable, $A_t k_t^\alpha$ and $A_{et} \theta k_{Rt}$ have to grow at the same rate. But balanced investment requires renewable k_{Rt} and production investment k_t to grow at the same rate, and both to grow at the rate of output g_y (since $F(\cdot)$ is constant returns to scale). Then

$$g_y = g_A + \alpha g_y = g_{A_e} + g_y$$

implies that $g_{A_e} = 0$. Put another way, when energy comes from accumulated capital, and as would be the case with firmed solar, total energy is linear in the amount of capital installed, stable growth in the long run requires that energy efficiency stop improving. All innovative resources are directed at capital and labor alone. This argument can be made rigorous in a decentralized world, which we do in a companion note.

Such a shift in the growth pattern could have quantitative bite. The estimates of Hassler et al. (2021) imply that for $g_{A_e} = 0$, the frontier for TFP growth g_A could rise from 1% to 1.2-1.4%, a significant improvement on what we have seen in recent decades. The possibilities of such an improvement are tantalizing. In Figure (16) we show the difference such an uptick makes in GDP per worker over the subsequent decades. By 2040, GDP per worker is 5-10.5% higher without the drag of finite resources for energy. While more speculative, this is more than double the gains from cheaper power studied above. It is worth emphasizing that the two macroeconomic effects of clean power are distinct. The first thing renewables do is make electricity cheaper in the medium run, an effect that is almost baked in at this point. The second thing they might do is free up innovative resources in the aggregate to better improve capital and labor productivity.

Figure 16: Removing the Resource Drag



Notes: This Figure shows real GDP per worker from 1947 to 2023 in 2017 chained dollars from the Bureau of Economic Analysis in black. The blue dashed line projects out the series to 2050 using the 1970-2023 average growth rate. The red dotted line uses the estimates from Hassler et al. (2021) to derive an estimate of additional growth that would result from replacing fossil fuels with renewable capital in the aggregate production function.

References

- Abhyankar, Nikit, Priyanka Mohanty, and Amol Phadke**, “Illustrative strategies for the United States to achieve 50% emissions reduction by 2030,” *Lawrence Berkeley National Laboratory*, 2021.
- Acemoglu, Daron**, “Directed Technical Change,” *The Review of Economic Studies*, 2002, 69 (4), 781–809.
- , **Philippe Aghion, Leonardo Bursztyn, and David Hemous**, “The Environment and Directed Technical Change,” *American Economic Review*, 2012, 102 (1), 131–166.
- , —, **Lint Barrage, and David Hémous**, “Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution,” Technical Report, National Bureau of Economic Research 2023.
- Arkolakis, Costas and Conor Walsh**, “Clean Growth,” Technical Report, National Bureau of Economic Research 2023.
- , **Arnaud Costinot, Dave Donaldson, and Andres Rodríguez-Clare**, “The Elusive Pro-Competitive Effects of Trade,” *mimeo*, 2012.
- Artuç, Erhan, Shubham Chaudhuri, and John McLaren**, “Trade Shocks and labor adjustment: A Structural Empirical Approach,” *American Economic Review*, 2010, 100 (3), 1008–1045.
- Bahn, Olivier, Mario Samano, and Paul Sarkis**, “Market power and renewables: The effects of ownership transfers,” *The Energy Journal*, 2021, 42 (4), 195–225.
- Baqae, David and Elisa Rubbo**, “Micro Propagation and Macro Aggregation,” *Annual Review of Economics*, 2023, 15 (1), 91–123.
- Barrage, Lint and William Nordhaus**, “Policies, projections, and the social cost of carbon: Results from the DICE-2023 model,” *Proceedings of the National Academy of Sciences*, 2024, 121 (13), e2312030121.
- Bartelme, Dominick, Arnaud Costinot, Dave Donaldson, and Andres Rodriguez-Clare**, “Sector-Level Economies of Scale: Estimation Using Trade Data,” 2017.
- Bilal, Adrien and Esteban Rossi-Hansberg**, “Anticipating Climate Change Across the United States,” Technical Report, National Bureau of Economic Research 2023.
- Bistline, John ET, Neil R Mehrotra, and Catherine Wolfram**, “Economic implications of the climate provisions of the inflation reduction act,” *Brookings Papers on Economic Activity*, 2023, 2023 (1), 77–182.

- Bistline, John, Geoffrey Blanford, Maxwell Brown, Dallas Burtraw, Maya Domeshek, Jamil Farbes, Allen Fawcett, Anne Hamilton, Jesse Jenkins, Ryan Jones et al.,** “Emissions and energy impacts of the Inflation Reduction Act,” *Science*, 2023, 380 (6652), 1324–1327.
- Cole, Wesley and A Will Frazier,** “Cost projections for utility-scale battery storage: 2023 update,” Technical Report, National Renewable Energy Lab.(NREL), Golden, CO (United States) 2023.
- Cruz, José-Luis and Esteban Rossi-Hansberg,** “The economic geography of global warming,” *Review of Economic Studies*, 2024, 91 (2), 899–939.
- Desmet, Klaus and Esteban Rossi-Hansberg,** “Climate change economics over time and space,” *Annual Review of Economics*, 2024, 16.
- , **Robert E Kopp, Scott A Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H Strauss,** “Evaluating the economic cost of coastal flooding,” Technical Report, National Bureau of Economic Research 2018.
- Dominish, Elsa, Nick Florin, and Sven Teske,** “Responsible Minerals Sourcing for Renewable Energy,” *Report prepared for Earthworks by the Institute for Sustainable Futures, University of Technology Sydney*, 2019, 54.
- Eckert, Fabian, Sharat Ganapati, and Conor Walsh,** “Urban-biased Growth: a Macroeconomic Analysis,” Technical Report, National Bureau of Economic Research 2022.
- Edenhofer, Ottmar, Lion Hirth, Brigitte Knopf, Michael Pahle, Steffen Schlömer, Eva Schmid, and Falko Ueckerdt,** “On the Economics of Renewable Energy Sources,” *Energy Economics*, 2013, 40, S12–S23.
- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J Klenow,** “How Destructive is Innovation?,” *Econometrica*, 2019, 87 (5), 1507–1541.
- Gervais, Antoine and J Bradford Jensen,** “The Tradability of Services: Geographic Concentration and Trade Costs,” *Journal of International Economics*, 2019, 118, 331–350.
- Gillingham, Kenneth, Richard G Newell, and William A Pizer,** “Modeling Endogenous Technological Change for Climate Policy Analysis,” *Energy Economics*, 2008, 30 (6), 2734–2753.
- Gonzales, Luis E, Koichiro Ito, and Mar Reguant,** “The Investment Effects of Market Integration: Evidence from Renewable Energy Expansion in Chile,” *Econometrica*, 2023, 91 (5), 1659–1693.

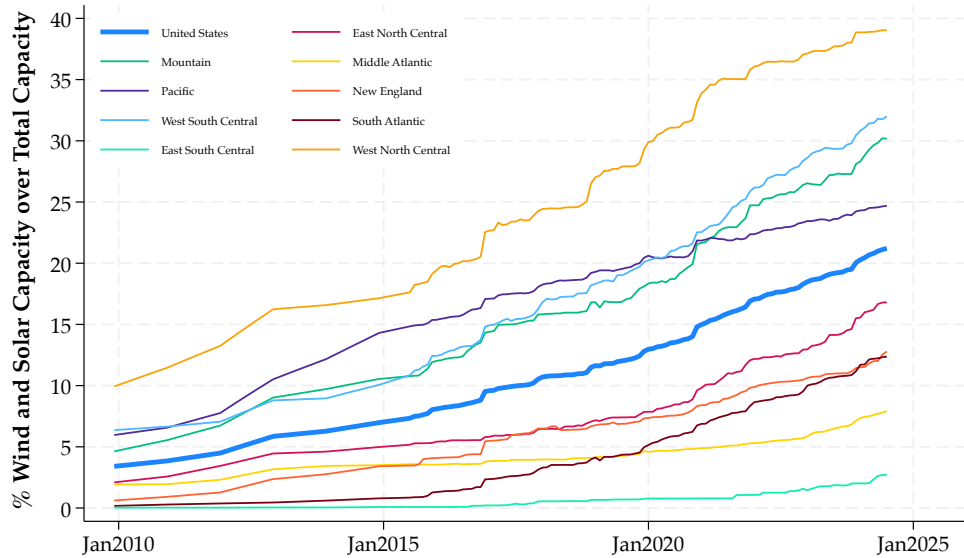
- Hall, Robert E**, “New Evidence on the Markup of Prices Over Marginal Costs and the Role of Mega-firms in the US Economy,” Technical Report, National Bureau of Economic Research 2018.
- Hanson, Gordon H**, “Local Labor Market Impacts of the Energy Transition: Prospects and Policies,” Technical Report, National Bureau of Economic Research 2023.
- Hassler, John, Per Krusell, and Conny Olovsson**, “Directed Technical Change as a Response to Natural Resource Scarcity,” *Journal of Political Economy*, 2021, 129 (11), 3039–3072.
- Hickel, Jason**, “What Does Degrowth Mean? A Few Points of Clarification,” *Globalizations*, 2021, 18 (7), 1105–1111.
- Hottman, Colin J, Stephen J Redding, and David E Weinstein**, “Quantifying the Sources of Firm Heterogeneity,” *The Quarterly Journal of Economics*, 2016, 131 (3), 1291–1364.
- Jenkins, Jesse D, Erin N Mayfield, Eric D Larson, Stephen W Pacala, and Chris Greig**, “Mission net-zero America: The nation-building path to a prosperous, net-zero emissions economy,” *Joule*, 2021, 5 (11), 2755–2761.
- , —, **Jamil Farbes, Ryan Jones, Neha Patankar, Qingyu Xu, and Greg Schivley**, “Preliminary report: The climate and energy impacts of the inflation reduction act of 2022,” *REPEAT Project*, 2022.
- Jones, Charles I and Dean Scrimgeour**, “The Steady-state Growth Theorem: Understanding Uzawa (1961),” *NBER Working Paper, August, (1961)*, 2005, pp. 1–9.
- Kanzig, Diego and Charles Williamson**, “The Drivers of Energy-saving Technical Change,” Technical Report, National Bureau of Economic Research 2024.
- Kypreos, Socrates and Olivier Bahn**, “A MERGE Model with Endogenous Technological Progress,” *Environmental Modeling & Assessment*, 2003, 8, 249–259.
- Labandeira, Xavier, José M Labeaga, and Xiral López-Otero**, “A meta-analysis on the price elasticity of energy demand,” *Energy policy*, 2017, 102, 549–568.
- Lemoine, Derek**, “Innovation-led Transitions in Energy Supply,” *American Economic Journal: Macroeconomics*, 2024, 16 (1), 29–65.
- Matsuyama, Kiminori and Philip Ushchev**, “Beyond CES: three alternative classes of flexible homothetic demand systems,” *Global Poverty Research Lab Working Paper*, 2017, (17-109).

- Mayfield, Erin, Jesse Jenkins, Eric Larson, and Chris Greig**, “Labor pathways to achieve net-zero emissions in the United States by mid-century,” *Energy Policy*, 2023, 177, 113516.
- Mehrotra, Neil**, “The Macroeconomics of Net Zero,” Technical Report, Minneapolis Federal Reserve 2024.
- Melitz, Marc J.**, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 2003, 71 (6), 1695–1725.
- Nolte, Christoph**, “High-resolution Land Value Maps Reveal Underestimation of Conservation Costs in the United States,” *Proceedings of the National Academy of Sciences*, 2020, 117 (47), 29577–29583.
- Nordhaus, William D**, “Optimal greenhouse-gas reductions and tax policy in the “DICE” model,” *The American Economic Review*, 1993, 83 (2), 313–317.
- , “Retrospective on the 1970s productivity slowdown,” 2004.
- , **Hendrik Houthakker**, and **Robert Solow**, “The Allocation of Energy Resources,” *Brookings Papers on Economic Activity*, 1973, 1973 (3), 529–576.
- , **Robert N Stavins**, and **Martin L Weitzman**, “Lethal Model 2: The limits to Growth Revisited,” *Brookings papers on economic activity*, 1992, 1992 (2), 1–59.
- Peters, Michael and Conor Walsh**, “Population Growth and Firm Dynamics,” Technical Report, National Bureau of Economic Research 2022.
- Pfenninger, Stefan, Adam Hawkes, and James Keirstead**, “Energy systems modeling for twenty-first century energy challenges,” *Renewable and Sustainable Energy Reviews*, 2014, 33, 74–86.
- Potter, Neal and Francis T Christy Jr**, “Trends in Natural Resource Commodities,” *Soil Science*, 1962, 94 (5), 349.
- Ramasamy, Vignesh, Jarett Zuboy, Michael Woodhouse, Eric O’Shaughnessy, David Feldman, Jal Desai, Andy Walker, Robert Margolis, and Paul Basore**, “US Solar Photovoltaic System and Energy Storage Cost Benchmarks, With Minimum Sustainable Price Analysis: Q1 2023,” Technical Report, National Renewable Energy Laboratory (NREL), Golden, CO (United States) 2023.
- Redding, Stephen J and Esteban Rossi-Hansberg**, “Quantitative Spatial Economics,” *Annual Review of Economics*, 2017, 9 (1), 21–58.
- Solow, RM**, “Intergenerational Equity and Exhaustible Resources,” *Review of Economic Studies*, 1974, 41 (5), 29–45.

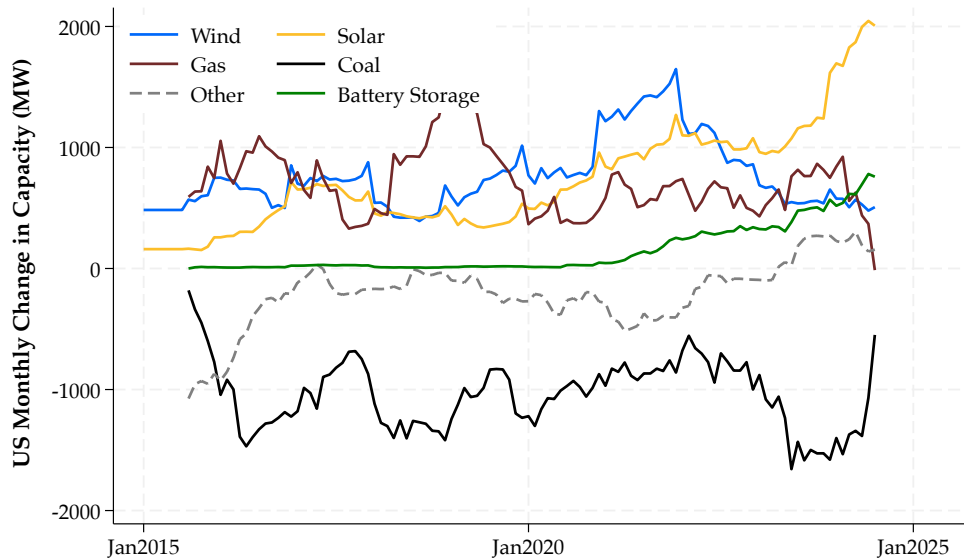
- Stehly, Tyler, Patrick Duffy, and Daniel Mulas Hernando**, “2022 Cost of wind energy review [Slides],” Technical Report, National Renewable Energy Laboratory (NREL), Golden, CO (United States) 2023.
- Stiglitz, Joseph E**, “A Neoclassical Analysis of the Economics of Natural Resources,” *Available at SSRN 250334*, 1980.
- Tong, Dan, David J Farnham, Lei Duan, Qiang Zhang, Nathan S Lewis, Ken Caldeira, and Steven J Davis**, “Geophysical Constraints on the Reliability of Solar and Wind Power Worldwide,” *Nature communications*, 2021, 12 (1), 6146.
- Uzawa, Hirofumi**, “Neutral Inventions and the Stability of Growth Equilibrium,” *The Review of Economic Studies*, 1961, 28 (2), 117–124.
- Way, Rupert, Matthew C Ives, Penny Mealy, and J Doyne Farmer**, “Empirically grounded technology forecasts and the energy transition,” *Joule*, 2022, 6 (9), 2057–2082.
- Ziegler, Micah S and Jessika E Trancik**, “Data series for Lithium-ion Battery Technologies,” <https://doi.org/10.7910/DVN/9FEJ7C>,, 2021.

A Figures and Tables

Figure 17: Renewable Investment Detail




(a) Solar and Wind Share in Nameplate Capacity



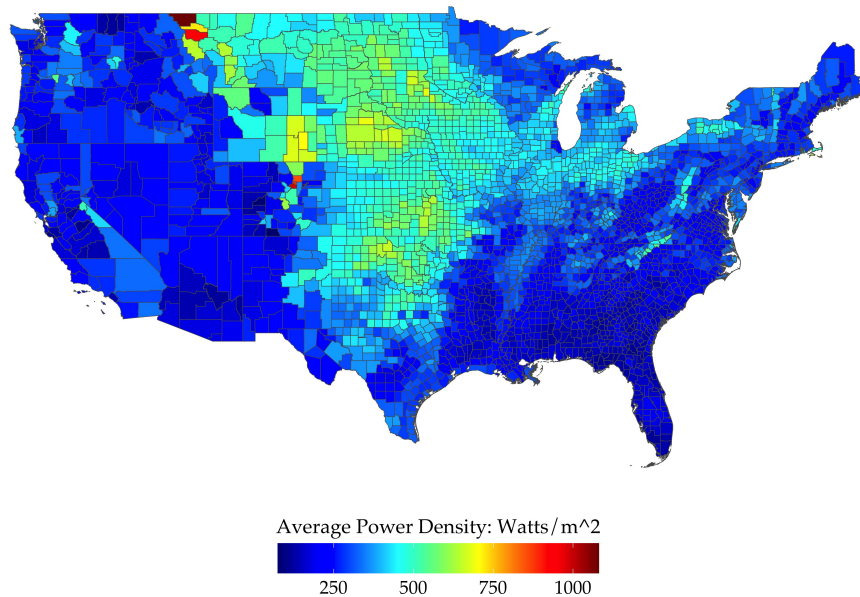
(b) Total Change In Capacity By Technology

Notes: Panel (a) of this figure shows the 12-month moving average of the monthly share of total electricity at the regional level coming from solar and wind. Panel (b) shows the change in nameplate capacity by technology. Data are from the US Energy Information Administration.

Figure 18: Solar and Wind Productivity Across Space



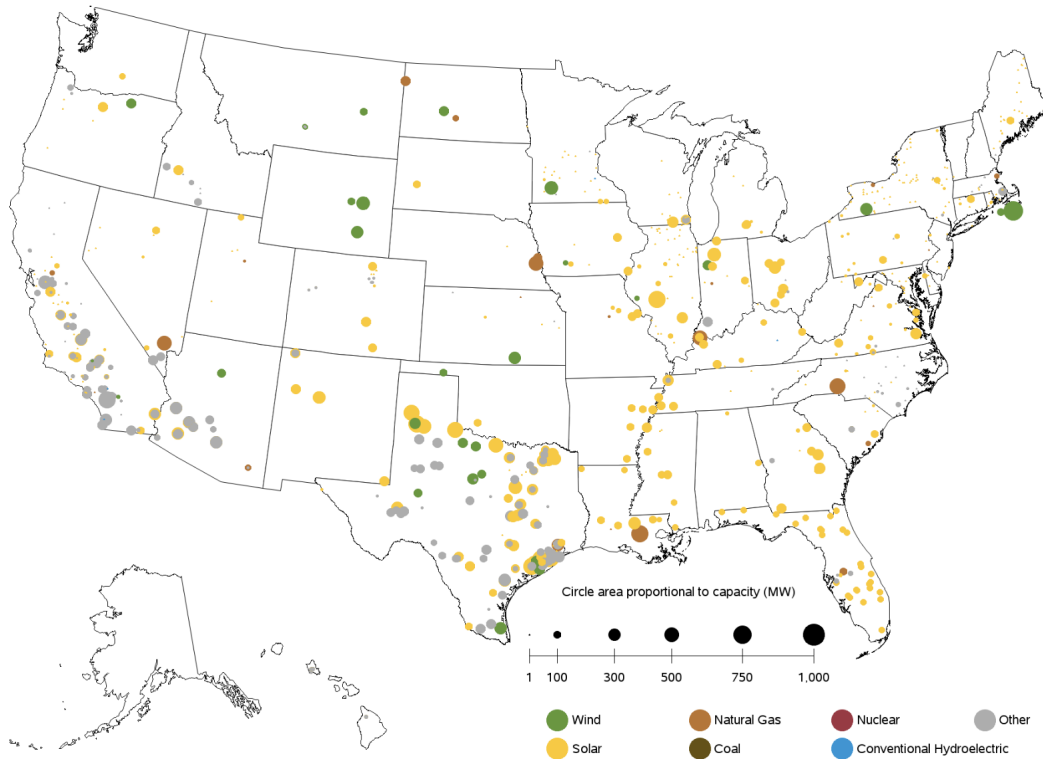
(a) Solar Potential



(b) Wind Potential

Notes: Panel (a) shows a measure of solar power potential, in average daily h produced by a 1 panel. Data is from the Global Solar Atlas. Panel (b) shows a measure of power output of a wind turbine, in average power density (watts per square meter), at a turbine height of 150 meters. Data is from the Global Wind Atlas. Units are US counties.

Figure 19: Energy Projects Currently Under Construction




Notes: Figure shows the distribution of energy generation plants by type coming online between July 2024 and June 2025, and is a reproduction of a graph by the Energy Information Administration (EIA). Other (in gray) mainly refers to lithium-ion storage plants.

Table 1: Top 25 Exposed Industries

Industry	Value for Φ_s^E / Φ_s^L
Alumina refining and primary aluminum production	0.88
Federal electric utilities	0.86
Paperboard mills	0.69
Industrial gas manufacturing	0.60
Cement manufacturing	0.49
Dairy cattle and milk production	0.47
Pulp mills	0.42
Iron, gold, silver, and other metal ore mining	0.41
Other real estate	0.39
Other basic organic chemical manufacturing	0.36
Nonferrous metal (except aluminum) smelting and refining	0.35
Ground or treated mineral and earth manufacturing	0.35
Other basic inorganic chemical manufacturing	0.32
Petroleum refineries	0.27
Plastics material and resin manufacturing	0.26
Iron and steel mills and ferroalloy manufacturing	0.24
Copper, nickel, lead, and zinc mining	0.21
Plastics bottle manufacturing	0.19
Asphalt paving mixture and block manufacturing	0.19
Paper mills	0.19
Fertilizer manufacturing	0.19
Gasoline stations	0.19
Lime and gypsum product manufacturing	0.18
Wet corn milling	0.18
Glass and glass product manufacturing	0.18

Notes: This Table reports the top 25 exposed industries to electricity price falls, as measured by Φ_s^E / Φ_s^L , the ratio of expenditure on electricity to total labor payments at the sectoral level. The data is from the Bureau of Economic Analysis Input-Output Tables for 2017 using the detailed 402 industry breakdown.

Figure 20: NREL Cell Efficiency By Type



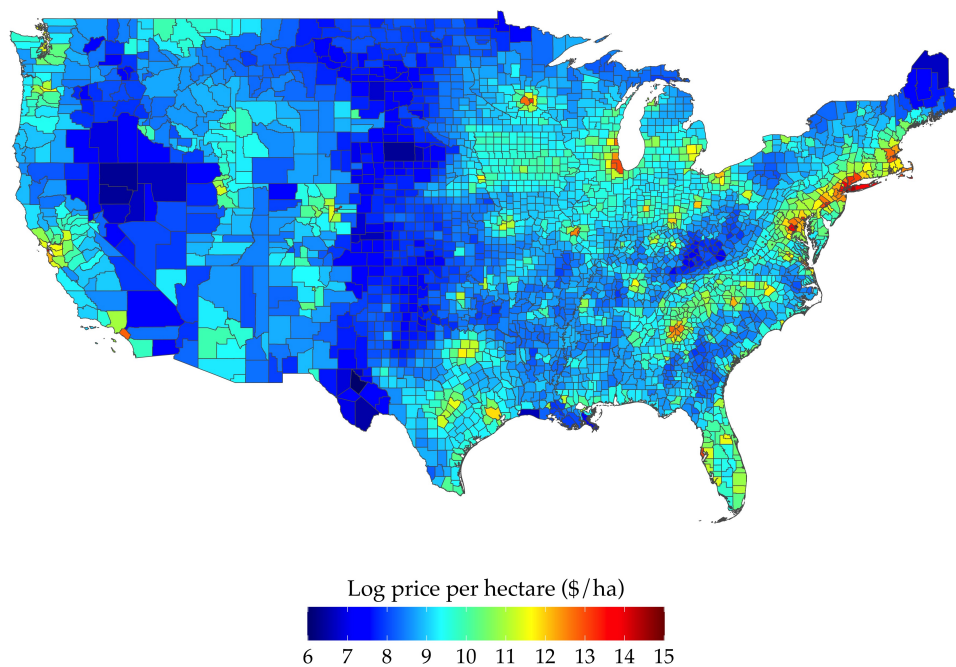
Notes: This Figure is a reproduction of a chart from the National Renewable Energy Laboratory, available here.

Table 2: Top 10 Large City Wage Increases

County	State	City	Wage Change (%)
San Bernardino	California	San Bernadino	4.01
Riverside	California	San Bernadino	3.84
Salt Lake	Utah	Salt Lake City	3.71
Orange	California	Los Angeles	3.5
Clark	Nevada	Los Vegas	3.46
Los Angeles	California	Los Angeles	3.29
King	Washington	Seattle	3.01
Santa Clara	California	Santa Clara	2.90
Tarrant	Texas	Dallas-Fort Worth	2.81
Bexar	Texas	San Antonio	2.78

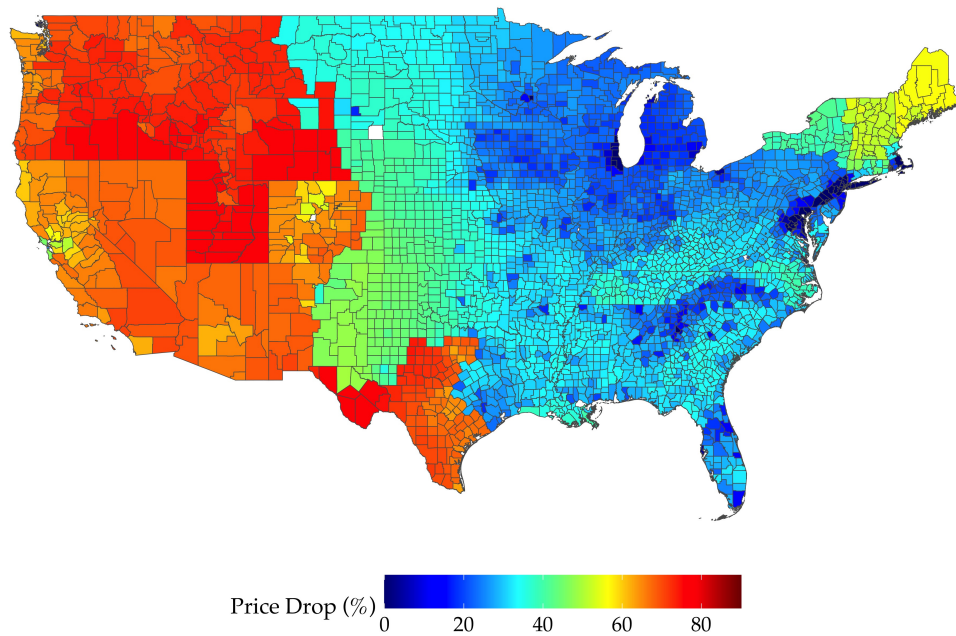
Notes: This Table reports the top 10 ten wage increases from the direct effect for counties with over 1 million employees.

Figure 21: Land Values in the US



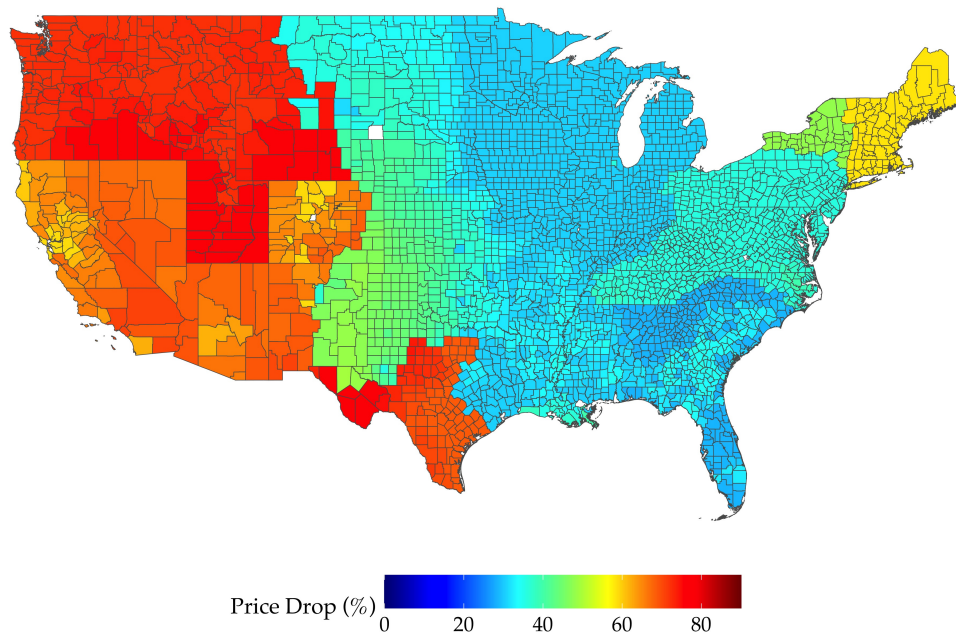
Notes: Figure plots estimates from Nolte (2020) for land prices in \$ per hectare. Estimates are plotted on a log scale. Values are averaged at the county level.

Figure 22: Implied Wholesale Price Drops between 2024 and 2040




Notes: Figure shows the implied bound fall in wholesale prices between 2024 and using equation (2). Current wholesale prices in 2024 are collected from the Energy Information Administration. We use the average at the RTO level for 10 price hubs: Northwest, ISO-NE, NYISO, ERCOT, Southwest, CAISO, PJM, MISO, SPP, FRCC and SERC. Units are US counties.

Figure 23: Implied Wholesale Price Drops with Mild Integration



Notes: Figure shows the implied bound fall in wholesale prices between 2024 and using equation (2), under the assumption that the maximum price is \$25 per MWh. Current wholesale prices in 2024 are collected from the Energy Information Administration. We use the average at the RTO level for 10 price hubs: Northwest, ISO-NE, NYISO, ERCOT, Southwest, CAISO, PJM, MISO, SPP, FRCC and SERC. Units are US counties.

Figure 24: Incorporating Labor Reallocation



Notes: Figure shows the calculated direct exposure measures Ω_l from equation (8) at the county level, against the sum of Ω_ℓ and the labor reallocation term in (8). The blue dashed line is the 45-degree line.

B A Model Of Trade, Power and Production

Consider a model of production and trade wherein agents live in a number of discrete locations ℓ . Suppose agents have preferences over an aggregator of sectoral goods given by

$$C = U(\{C_s\}_s). \quad (11)$$

Workers have preferences over the final consumption aggregator and residential housing h . The price of this aggregator is chosen as the numeraire. Each period, an individual worker j of labor type chooses a location ℓ , sector s , and quantities of housing (h) and the final good (c) to solve the following utility-maximization problem:

$$\max_{\ell} \{ \vartheta_{\ell}^j \mathbb{E}_{\vartheta_s} \max_{s,h,c} \{ V^W(c,h) \vartheta_s^j \} \} \quad \text{subject to} \quad m_{\ell}^h h + c = w_{\ell s}, \quad (12)$$

where $V^W(c,h)$ is concave and continuously differentiable, and m_{ℓ}^h is the rental rate on residential land in location ℓ . ϑ_{ℓ}^j and ϑ_s^j are sectoral preference shocks that give rise to smooth labor supply curves, discussed further in Section (B.5). Time is discrete, but for the moment we suppress the time index t .

B.1 Firms

Firms produce a unique differentiated variety i . Firm i located in location ℓ and sector s produce with

$$y_i = z_i F_{\ell s}(l, e, \mathbf{x})$$

where l is labor, e is electricity and \mathbf{x} is a vector of intermediate inputs. z_i is an index of firm level TFP. Note that the production function $F_{\ell s}(\cdot)$ may be location- and sector-specific, incorporating exogenous local productivity differences and endogenous agglomeration forces, as long as these are taken as given by the firm.

Firms need to pay an entry cost, defined by $g_{\ell s}(l, e) = 1$, and after entry they draw their productivity from an exogenous distribution $\Psi_{\ell s}(z)$, which may depend on location and sector. They exit at constant rate ζ . We assumed g is constant returns to scale. The resulting entry cost is denoted $G_{\ell s}(w_{\ell s}, p_{\ell}^{\mathcal{E}})$.

We suppose there are no trade costs, and that intermediate inputs enter in a single aggregate input for all firms in the same way (though they may use this with different intensity).

That is, we suppose

$$y_i = z_i F_{\ell s}(l, e, X),$$

where

$$X = f(x)$$

is a symmetric aggregator for all firms. We further assume it takes the same form as final goods aggregation, so that both the intermediates price and the final good price serve as the numeraire. Cost minimization for a given level of output y allows us to write a cost function

$$C_{\ell s}(y; z) = z^{-1} y v_{\ell s}(w_{\ell s}, p_{\ell}^{\mathcal{E}}, y),$$

where v is the average unit cost function.

The pricing decision leads to a concave revenue function $R_s(y) = D_s r_s(y)$, where D_s is an aggregate sectoral demand shifter. Notice that this is not an innocuous assumption. For example, in our context it implies the absence of trade costs or differential wedges across sectors. Nevertheless, it does allow for a broad class of demand functions. Consider for example single-aggregator demand functions such as those considered by Arkolakis et al., 2012; Matsuyama and Ushchev, 2017,

$$y_s(p) = d_s \left(\frac{p}{P_s^*} \right),$$

where P_s^* is an aggregator function, p is the firm's price and $d_s(\cdot)$ is a demand function that is strictly decreasing in its argument. The revenue function can then be written as

$$R_s(y) = py = P_s^* d_s^{-1}(y) y = d_s^{-1}(y) y \times P_s^* = r_s(y) \times D_s.$$

where $D_s \equiv P_s^*$ and $r_s(y) \equiv d_s^{-1}(y) y$. The CES demand is a special case of this class with

$$y_s(p_s) = p_s^{-\sigma_s} \times \frac{Y}{P_s^{1-\sigma_s}},$$

where Y is total income and P_s is the CES aggregator.

B.2 Capitalists

There is a population of capitalists who own the firms, local land, solar capital and intermediates goods. They care only about consumption C_t^K , with the same aggregation over

sectors as in (11), and have intertemporal preferences given by

$$V^K = \sum_{t=0}^{\infty} \beta^t v(C_t^K). \quad (13)$$

We suppose that they can invest in new firms, and the vector of intermediates which depreciate at rate δ_s and have rental rate m_{st} . The stock of each intermediate is denoted X_{st} . They receive the profits from all the firms and land. Their budget constraint is

$$C_t^K + \sum_{\ell} G_{\ell s}(w_{\ell st}, p_{\ell t}^{\mathcal{E}})(N_{\ell st+1} - (1 - \xi)N_{\ell st}) + \sum_{\ell} Q_{t+1}(S_{\ell t+1} - (1 - \delta^S)S_{\ell t}) = \sum_s \sum_l \Pi_{s\ell t} + \sum_l p_{\ell t}^{\mathcal{E}} \theta_{\ell} S_{\ell t+1} + \sum m_{\ell t}^h H, \quad (14)$$

where $N_{\ell st}$ is the number of firms in location ℓ and sector s at time t , $\Pi_{\ell st}$ is the profits they make, and H is the supply of residential land (in fixed supply).

Free entry into firm creation requires that the return on creating firms is equal to the return on investing in the intermediates, so that

$$G_{\ell s}(w_{\ell st}, p_{\ell t}^{\mathcal{E}}) = \sum_{\tau=0}^{\infty} R_{t \rightarrow t+\tau} (1 - \xi)^{\tau} \left(\int_z \left[\max_y D_{st+\tau} r(y) - C_{\ell st+\tau}(y; w_{\ell st+\tau}, p_{\ell t+\tau}^{\mathcal{E}}, z) \right] d\Psi_{\ell s}(z) \right). \quad (15)$$

where $R_{t \rightarrow t+\tau}$ is the common real cumulative return on assets (which must be common for all assets in equilibrium).

Solar capital $S_{\ell t}$ provides θ_{ℓ} units of electricity per period. It can be bought by converting units of the final good into capital at rate $1/Q_t$, which we take to be an exogenous parameter. It depreciates at rate δ^S .

B.3 Definition of Equilibrium

An equilibrium is a time path for of wages $w_{\ell st}$ in each location-sector, a path for sectoral intermediates prices P_{st} , rental rates on residential land $\{m_{\ell t}^h\}$, an allocation of consumption C_t^K for capitalists, stocks of workers in each location-sector $L_{\ell st}$, numbers of firms in each location $N_{\ell st}$, stocks of solar capital $S_{\ell t}$, such that given a price for solar capital Q_t ,

1. Workers solve their problem (12) statically,
2. Capitalists maximize (13) subject to (14),
3. The free entry condition (15) holds,

4. All markets clear.

A steady-state equilibrium is one in which all prices and allocations are constant through time.

B.4 General Equilibrium Price Changes

Proof of Proposition 1. In a steady state of the model with no growth, the free-entry condition can be written

$$\begin{aligned} G_{\ell s}(w_{\ell s}, p_{\ell}^{\varepsilon}) &= \kappa \int \pi_{\ell s}(z) d\Psi_{\ell s}(z) \\ &= \kappa \int \max_y \left[D_s r_s(y) - z^{-1} y v(w_{\ell s}, p_{\ell}^{\varepsilon}, y) \right] d\Psi_{\ell s}(z), \end{aligned}$$

where κ is a proportional constant equal to $(\beta + 1) / (\beta \xi + 1)$, and $G_{\ell s}(w_{\ell s}, p_{\ell}^{\varepsilon})$ is the optimized entry cost.

We derive the first order response of local factor prices to a decrease in electricity prices as follows. By the envelope theorem, we have

$$\frac{\partial \pi_{\ell s}(z, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial w_{\ell s}} = -z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial w_{\ell s}}, \quad \frac{\partial \pi_{\ell s}(z, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial p_{\ell}^{\varepsilon}} = -z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial p_{\ell}^{\varepsilon}},$$

where y is understood to be optimal output at the given vector of factor prices. In addition, when sectoral demand changes the effect on profit is given by

$$\frac{\partial \pi_{\ell s}(z)}{\partial D_s} = r_s(y).$$

Totally differentiating the free-entry condition and using these expressions yields

$$\begin{aligned} \frac{\partial G_{\ell s}}{\partial w_{\ell s}} dw_{\ell s} + \frac{\partial G_{\ell s}}{\partial p_{\ell}^{\varepsilon}} dp_{\ell}^{\varepsilon} &= \kappa \int \left[r_s(y) D_s d \log D_s \right. \\ &\quad \left. - z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial w_{\ell s}} dw_{\ell s} - z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial p_{\ell}^{\varepsilon}} dp_{\ell}^{\varepsilon} \right] d\Psi_{\ell s}(z). \end{aligned} \tag{16}$$

The intermediate bundle is the numeraire and so receives no price change. We can also

write the free-entry condition using Euler's theorem and Shephard's Lemma as

$$\frac{\partial G_{\ell s}}{\partial p_{\ell}^{\varepsilon}} w_{\ell s} + \frac{\partial G_{\ell s}}{\partial p_{\ell}^{\varepsilon}} p_{\ell}^{\varepsilon} = \kappa \int \left[r_s(y) D_s - z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial w_{\ell s}} w_{\ell s} - z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial p_{\ell}^{\varepsilon}} p_{\ell}^{\varepsilon} - z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial P} \right] d\Psi_{\ell s}(z), \quad (17)$$

where P is the price of the intermediate bundle/final good, and is the numeraire so equals 1. Use this last expression to rearrange and obtain

$$\int r_s(y) D_s d\Psi_{\ell s}(z) = \frac{\partial G_{\ell s}}{\partial w_{\ell s}} w_{\ell s} + \frac{\partial G_{\ell s}}{\partial p_{\ell}^{\varepsilon}} p_{\ell}^{\varepsilon} + \kappa \int \left[z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial w_{\ell s}} w_{\ell s} + z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial p_{\ell}^{\varepsilon}} p_{\ell}^{\varepsilon} + z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial P} \right] d\Psi_{\ell s}(z),$$

and insert this into (16) to get

$$\begin{aligned} & \left[\frac{\partial G_{\ell s}}{\partial w_{\ell s}} w_{\ell s} + \kappa \int z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial w_{\ell s}} w_{\ell s} d\Psi_{\ell s}(z) \right] d \log w_{\ell s} + \left[\frac{\partial G_{\ell s}}{\partial p_{\ell}^{\varepsilon}} p_{\ell}^{\varepsilon} + \kappa \int z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial p_{\ell}^{\varepsilon}} p_{\ell}^{\varepsilon} d\Psi_{\ell s}(z) \right] d \log p_{\ell}^{\varepsilon} \\ &= \left[\frac{\partial G_{\ell s}}{\partial w_{\ell s}} w_{\ell s} + \kappa \int z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial w_{\ell s}} w_{\ell s} d\Psi_{\ell s}(z) + \frac{\partial G_{\ell s}}{\partial p_{\ell}^{\varepsilon}} p_{\ell}^{\varepsilon} + \kappa \int z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial p_{\ell}^{\varepsilon}} p_{\ell}^{\varepsilon} d\Psi_{\ell s}(z) + \int z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial P} d\Psi_{\ell s}(z) \right] d \log D_s. \end{aligned}$$

Now note that total expenditure on labor (inclusive of both entry costs and variable costs) is

$$\Phi_{\ell s}^L \equiv N_{\ell s} \left[\frac{\partial G_{\ell s}}{\partial w_{\ell s}} w_{\ell s} + \kappa \int z^{-1} y \frac{\partial v_{\ell s}(y, w_{\ell s}, p_{\ell}^{\varepsilon})}{\partial w_{\ell s}} w_{\ell s} d\Psi_{\ell s}(z) \right],$$

with a similar expression for $\Phi_{\ell s}^E$ and $\Phi_{\ell s}^X$, where $N_{\ell s}$ is the number of firms in location ℓ and sector s . Rearranging gives us

$$d \log w_{\ell s} = -\frac{\Phi_{\ell s}^E}{\Phi_{\ell s}^L} d \log p_{\ell}^{\varepsilon} + \frac{\Phi_{\ell s}^E + \Phi_{\ell s}^L + \Phi_{\ell s}^X}{\Phi_{\ell s}^L} d \log D_s,$$

where $\Phi_{\ell s}^E$ is total local sectoral expenditure on electricity, $\Phi_{\ell s}^L$ is expenditure on labor and $\Phi_{\ell s}^X$ is local expenditure on intermediates.

B.5 Labor Supply Elasticities Across Sectors

Here we put structure on worker preferences, and derive their resulting labor supply curves. We first assume that workers' preferences over consumption and housing are Cobb-Douglas, and the weight on housing is α .

We assume that workers draw their idiosyncratic preference shocks for each location and sector from Frechet distributions with scale parameters A_ℓ and $B_{\ell s}$ respectively (for location, and sector conditional on location), and shape parameters ϱ and η . The fraction of workers deciding to live in location ℓ, λ_ℓ , and for working in sector s conditional on choosing to live in location $\ell, \mu_{\ell s}$, are given by

$$\lambda_\ell = \frac{A_\ell (r_\ell^{-\alpha} \Psi_\ell)^\varrho}{\sum_{\ell'} A_{\ell'} (r_{\ell'}^{-\alpha} \Psi_{\ell'})^\varrho} \quad \mu_{\ell s} = \frac{B_{\ell s} (w_{\ell s})^\eta}{\sum_{s'} B_{\ell s'} (w_{\ell s'})^\eta},$$

where $\Psi_\ell \equiv (\sum_s B_{\ell s} (w_{\ell s})^\eta)^\frac{1}{\eta}$. From this expression, we can derive the change in sectoral employment shares as

$$d \log \mu_{\ell s} = \eta d \log w_{\ell s} - \eta \sum \mu_{\ell s} d \log w_{\ell s}.$$

Then we have

$$\sum_s \frac{\mu_{\ell s} w_{\ell s}}{\sum_s \mu_{\ell s} w_{\ell s}} d \log \mu_{\ell s} = \eta \sum_s \frac{\mu_{\ell s} w_{\ell s}}{\sum_s \mu_{\ell s} w_{\ell s}} d \log w_{\ell s} - \eta \sum_s \mu_{\ell s} d \log w_{\ell s}.$$

B.6 General Equilibrium Aggregate Demand Effects

We move to consider the aggregate demand effects in the model above by making some further parametric assumptions. Assume

1. There are no trade costs
2. There is no intermediate usage

$$F_{\ell s}(l, e, X) = F_{\ell s}(l, e)$$

3. The utility function for final demand is Cobb-Douglas, so that

$$U(\{C_s\}_s) = \prod_{s=1}^S C_s^{\gamma_s}$$

4. Aggregation within sectors is CES with elasticity of substitution σ_s
5. The entry cost is denominated in units of the final good, so that the cost of entry is equal to \bar{g} and the same everywhere

6. Production functions are constant returns to scale

In this case it can be shown that the sectoral demand shifter on firm profits is given by

$$D_s = \frac{\gamma_s Y}{P_s^{1-\sigma_s}},$$

where Y is aggregate income, given by

$$Y = \sum_{\ell} \sum_s w_{\ell s} L_{\ell s} + \sum_{\ell} \sum_s p_{\ell}^{\varepsilon} E_{\ell s} + \sum_{\ell} \sum_s N_{\ell s} \int \pi_{\ell s}(z) d\Psi_{\ell s}(z), \quad (18)$$

or the sum of labor income, electricity sales and firm profits. For notational convenience, call aggregate profits $\Pi \equiv \sum_{\ell} \sum_s N_{\ell s} \int \pi_{\ell s}(z) d\Psi_{\ell s}(z)$. In turn, the price index is given by

$$P_s^{1-\sigma_s} = \sum_{\ell} N_{\ell s} \int p_{\ell s}(z)^{1-\sigma} d\Psi_{\ell s}(z),$$

where $p_{\ell s}(z)$ is the intermediate good price for a firm with productivity z in location ℓ and sector s . As such,

$$d \log D_s = d \log Y - d \log P_s^{1-\sigma_s},$$

with the additional restriction that

$$\sum \gamma_s d \log P_s = 0,$$

by choice of the numeraire. Consider first the change in the sectoral price index.

$$\begin{aligned} d \log P_s^{1-\sigma_s} &= \frac{\sum_{\ell} N_{\ell s} \int p_{\ell s}(z)^{1-\sigma} d\Psi_{\ell s}(z) d \log N_{\ell s}}{\sum_{\ell} N_{\ell s} \int p_{\ell s}(z)^{1-\sigma} d\Psi_{\ell s}(z)} \\ &\quad + (1-\sigma) \frac{\sum_{\ell} N_{\ell s} \int p_{\ell s}(z)^{1-\sigma} \left(\frac{w_{\ell s} l_{\ell s}(z)}{w_{\ell} l_{\ell s}(z) + p_{\ell}^{\varepsilon} e_{\ell s}(z)} d \log w_{\ell s} + \frac{p_{\ell}^{\varepsilon} e_{\ell s}(z)}{w_{\ell} l_{\ell s}(z) + p_{\ell}^{\varepsilon} e_{\ell s}(z)} d \log p_{\ell}^{\varepsilon} \right) d\Psi_{\ell s}(z)}{\sum_{\ell} N_{\ell s} \int p_{\ell s}(z)^{1-\sigma} d\Psi_{\ell s}(z)}. \end{aligned}$$

where the second term invokes Shepard's Lemma. If production functions are constant returns to scale, then cost shares do not vary with firm level efficiency z , and this can be

written as

$$d \log P_s^{1-\sigma_s} = \frac{\sum_{\ell} N_{\ell s} \int p_{\ell s}(z)^{1-\sigma} d\Psi_{\ell s}(z) d \log N_{\ell s}}{\sum_{\ell} N_{\ell s} \int p_{\ell s}(z)^{1-\sigma} d\Psi_{\ell s}(z)} + (1-\sigma) \frac{\sum_{\ell} N_{\ell s} \int p_{\ell s}(z)^{1-\sigma} d\Psi_{\ell s}(z) \left(\frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} d \log w_{\ell s} + \frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} d \log p_{\ell}^{\varepsilon} \right)}{\sum_{\ell'} N_{\ell' s} \int p_{\ell' s}(z)^{1-\sigma} d\Psi_{\ell' s}(z)}.$$

Note that market clearing in the output market has

$$\frac{\sigma}{\sigma-1} (w_{\ell s} L_{\ell s} + p_{\ell}^{\varepsilon} E_{\ell s}) = \frac{N_{\ell s} \int p_{\ell s}(z)^{1-\sigma} d\Psi_{\ell s}(z)}{\sum_{\ell'} N_{\ell' s} \int p_{\ell' s}(z)^{1-\sigma} d\Psi_{\ell' s}(z)} \gamma_s Y.$$

So we can rewrite this as

$$d \log P_s^{1-\sigma_s} = \frac{\sigma}{\sigma-1} \sum_{\ell} \frac{w_{\ell s} L_{\ell s} + p_{\ell}^{\varepsilon} E_{\ell s}}{\gamma_s Y} d \log N_{\ell s} - \frac{1}{\sigma} \sum_{\ell} \frac{w_{\ell s} L_{\ell s} + p_{\ell}^{\varepsilon} E_{\ell s}}{\gamma_s Y} \left(\frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} d \log w_{\ell s} + \frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} d \log p_{\ell}^{\varepsilon} \right).$$

Now the free entry condition requires expected profit to equal the entry cost \bar{g} , so that

$$\kappa \int_z (p_{\ell s}(z) - v_{\ell s}(w_{\ell s}, p_{\ell}^{\varepsilon})) y_{\ell s}(z) d\Psi_{\ell s} = \bar{g}$$

implies

$$\frac{1}{\sigma} \mathbb{E}_{\ell s}[\text{sales}] = \bar{g} / \kappa,$$

given the optimal pricing formula with CES demand. In addition, total sales equaling factor income and profits requires

$$N_{\ell s} \mathbb{E}_{\ell s}[\text{sales}] = \frac{\sigma}{\sigma-1} (w_{\ell s} L_{\ell s} + p_{\ell}^{\varepsilon} E_{\ell s}),$$

so that

$$N_{\ell s} = \frac{1}{(\sigma-1)} \kappa \bar{g}^{-1} (w_{\ell s} L_{\ell s} + p_{\ell}^{\varepsilon} E_{\ell s}).$$

As such

$$\begin{aligned} d \log N_{\ell s} &= d \log \left(w_{\ell s} L_{\ell s} + p_{\ell}^{\mathcal{E}} E \right). \\ &= \left(\frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} (d \log w_{\ell s} + d \log L_{\ell s}) + \frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} (d \log p_{\ell}^{\mathcal{E}} + d \log E_{\ell s}) \right) \end{aligned}$$

Similarly, with profits being a fixed fraction of revenue under the CES demand structure, we also have from equation (18)

$$\begin{aligned} d \log Y &= \sum_{\ell} \sum_s \frac{\sigma / (\sigma - 1) (w_{\ell s} L_{\ell s} + \bar{p}_l E)}{Y} \left(\frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} (d \log w_{\ell s} + d \log L_{\ell s}) \right. \\ &\quad \left. + \frac{\Phi_{\ell s}^L}{\Phi_{\ell s}^L + \Phi_{\ell s}^E} (d \log p_{\ell}^{\mathcal{E}} + d \log E_{\ell s}) \right). \end{aligned}$$

C Price Bounds for Wind Power

We conduct our analysis in the main text using estimates for firmed solar power, as we believe it is likely to be the dominant technology in a renewable dominated grid, owing both to its rapid cost falls and a relatively unconstrained local potential in the analysis above. However, wind power has a significantly different geography of potential (see Figure 18). In particular, wind power is strong in some places where solar is weak; in particular in the rustbelt of Ohio, Michigan, and upstate New York. It is worth asking how our estimates of price drops would change if we extend the analysis to consider wind power.

Wind power is somewhat more involved to model than solar. The exact capacity factors depend on the turbine model used, which are currently quite heterogeneous, and difficult to project out to the future. As such, we take the following approach. We gather data for current capacity factors of different turbine classes and wind power density at 100m from the Global Wind Atlas. We use this data to form an estimate of wind power output at the local level.

We then assume that the capital cost of turbines continues to track the rate of decline seen in the Levelized Cost of Energy database for major countries from the International Renewable Energy Database, at -4% a year, and project this out to 2040. In addition, we assume that land costs remain the same, currently accounting for 2% of the levelized cost of wind in the United States using data from the NREL (Stehly et al., 2023). As we do for

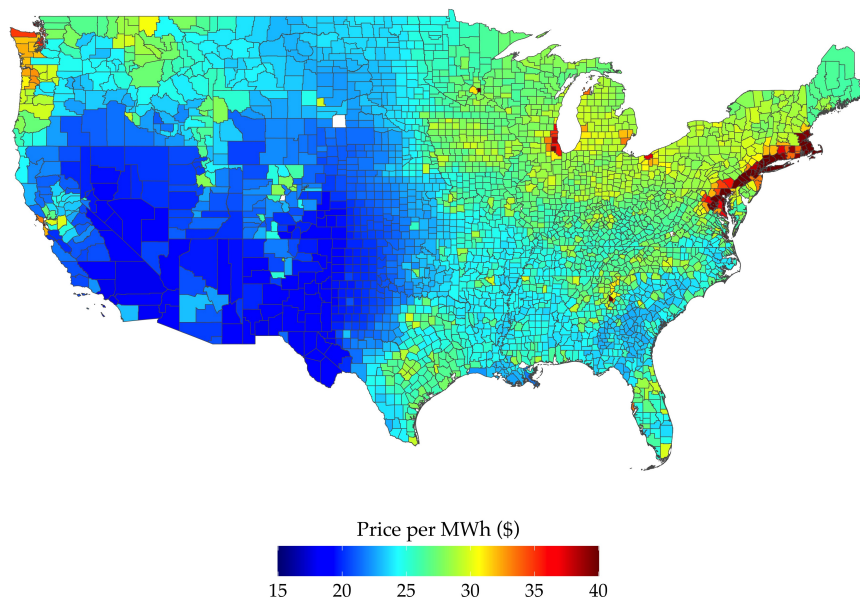
solar, we scale land costs depending on local land prices, and assume a long-run cost of financing of 5%. Lastly, we assume a thirty year life span, and a depreciation rate of 2%.

Together, this allows us to develop the price bound in (2) for wind. The results are plotted in Figure 25. Wind turns out to imply a lower cost bound than solar in select locations. In particular, the Rustbelt as a whole sees lower prices coming from wind than solar. However, for much of the country it is the solar bound that dominates, which can be seen by taking the minimum of the two bounds and plotting them in Figure 26.

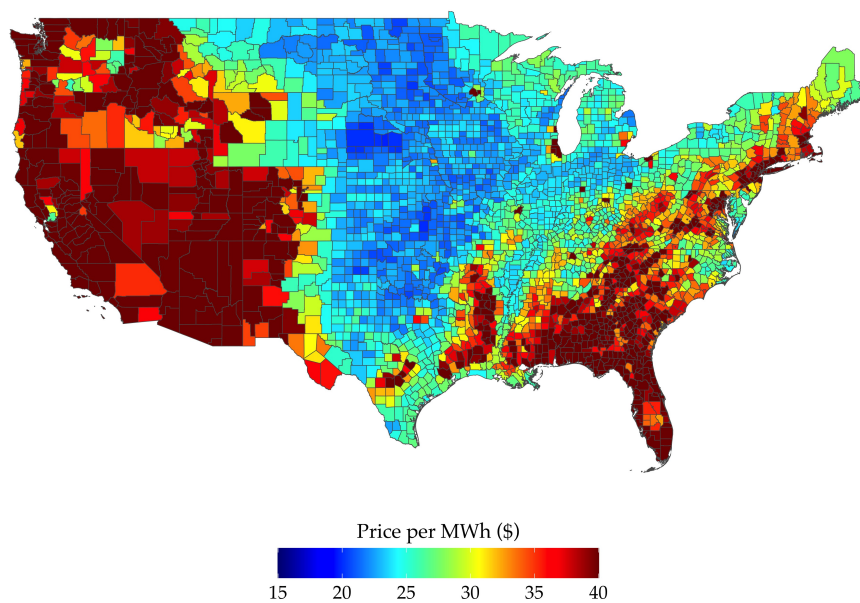
For this exercise we do not add storage costs to wind, as we do for solar, as the intermittency cost of wind energy is somewhat less pronounced, and likely to be further mitigated by the high amounts of storage required for solar in a renewable-dominated grid. Adding the cost of 8 hours of storage assumed above to the capital cost of wind in 2040 raises the long run price bounds below by around 20% across the US.

We are hesitant to include this analysis in the main body for a simple reason: land constraints for wind are far more severe than for solar. The NREL estimates that total developable capacity for wind is less than a tenth of the 83 terrawatts of solar potential. The land requirement per megawatt is also an order of magnitude greater, with many of the best sites in the US already developed. This makes us doubt the applicability of the free entry condition in every county, whereas for firmed solar we feel this is a much more reasonable assumption.

Figure 25: Price Maps Using Solar and Wind Technology



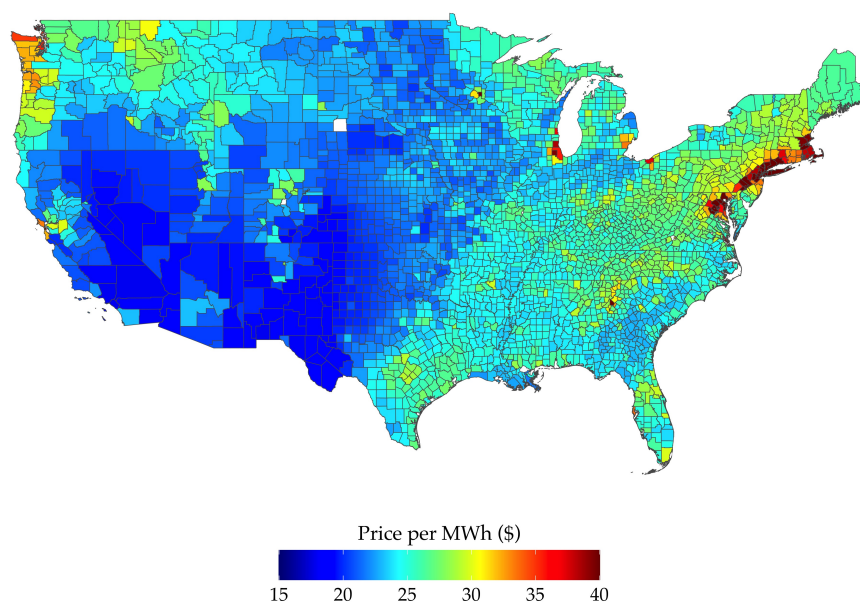
(a) Solar



(b) Wind

Notes: Panel (a) shows the price bound from equation (2) using data for firmed solar, and Panel (b) does the same using data for wind energy.

Figure 26: Minimum Of Solar and Wind Price Bound



Notes: This Figure presents the minimum implied price bound at the local level from the subfigures of Figure (25).