45V EXEMPTIONS NEED STRONG GUARDRAILS TO PROTECT CLIMATE, GROW HYDROGEN INDUSTRY

DAN ESPOSITO, ERIC GIMON, AND MIKE O'BOYLE

FEBRUARY 2024[i]

EXECUTIVE SUMMARY

This paper summarizes new analysis and recommendations in response to a December 2023 Notice of Public Rulemaking (NPRM) from the United States Treasury Department and the Internal Revenue Service (hereafter "Treasury"), which proposes regulations for the Inflation Reduction Act's (IRA) Section 45V Clean Hydrogen Production Tax Credit (45V).[1]

Treasury's draft guidance governs how electrolyzers—which split hydrogen from water using electricity—must source power to ensure a low greenhouse gas (GHG) emissions impact and earn the lucrative tax credit. While Treasury's core framework is strong, it considers several exemptions, including an indiscriminate carve-out for some percentage of existing clean energy to qualify as having no emissions impact.

Our analysis shows exempting just 5 percent of existing clean energy from Treasury's requirements would allow for approximately 1.5 million metric tons (MMT) of annual dirty hydrogen production, contributing roughly 30 to 60 MMT of CO_2 per year (equivalent to putting 13.3 million gasoline cars on the road) by inducing fossil fuel power plants to run more often. Such a broad exemption would also harm the hydrogen industry's long-term growth by eroding the incentive to invest in flexible electrolyzers capable of sustained success after 45V expires. Other, more targeted exemptions could be similarly disastrous for climate and the U.S. hydrogen industry's competitiveness if Treasury doesn't take extreme care.

ENERGY INNOVATION
POLICY & TECHNOLOGY LLC®

www.energyinnovation.org

98 Battery Street, Suite 202
San Francisco, CA 94111
policy@energyinnovation.org

In April 2023, we shared our analysis and recommendations for smart design of the 45V credit to reduce GHG emissions and successfully grow the clean hydrogen industry over the long term.[2] We recommended Treasury adopt strict guardrails in the form of the "three pillars" of incrementality, deliverability, and hourly time-matching to ensure electrolyzers earning the tax credit would have to source power in a manner that would achieve the upstream and production GHG emissions requirements mandated by Congress in the IRA.[ii] In its NPRM, Treasury proposes to adopt these requirements but seeks comment on potential exemptions.

In this paper, we assess the components of Treasury's NPRM related to the "incrementality" requirement as applied to electrolyzers seeking to earn the tax credit. Specifically, we analyze Treasury's proposal for a "general carve-out" that would exempt 5 percent of existing U.S. clean energy generation from the incrementality requirement.[iii] We also provide recommendations related to whether and how to offer targeted exemptions for the use of clean energy that might otherwise be curtailed, retire, or have its emissions impacts offset elsewhere.

Key takeaways

- **Protecting Treasury's draft rules around incrementality, deliverability, and hourly time-matching is essential to preserving the emissions benefits of 45V and promoting clean hydrogen's long-term growth.** The proposed rules strike the appropriate balance between accurately verifying compliance with 45V's GHG emissions thresholds and being readily administrable.
    - Should Treasury offer any exemptions, they should not sacrifice similar rigor in ensuring a high level of accuracy in determining GHG emissions intensity for ease of earning the credit, as doing so would undermine the entire program's integrity.
    - In particular, any error that inadvertently exempts an electrolyzer's use of existing clean energy in a manner that actually induces a net increase in GHG emissions would subsidize hydrogen production with a GHG emissions intensity of up to 50 to 100 times 45V's statutorily-required threshold.
- **A "general carve-out" that would allow any percentage of existing U.S. clean energy generation to be unconditionally exempt from the incrementality requirement would undermine the emissions and industry growth objectives of the IRA.** This approach would do a very poor job at meeting its intent of qualifying existing clean power that, if used, would not induce GHG emissions elsewhere on the grid.

[ii] Our paper used the term "additionality" for what Treasury calls "incrementality" in its NPRM.
[iii] Treasury proposes a 5 percent carve-out but asks whether this would be the appropriate number. It also requests comment on whether a higher amount, such as 10 percent, would be appropriate.

- Specifically, even a 5 percent exemption would allow for approximately 1.5 MMT of dirty electrolytic hydrogen production per year, contributing roughly 30 to 60 MMT of CO_2 emissions annually—with a 10 percent exemption doubling this negative outcome.
- Any indiscriminate exemption for existing clean energy would also cause much more damage to the extent it hampers innovation. For example, electrolyzers might be able to make hydrogen in most hours by drawing electricity from new wind and solar projects, but by doing so, they must still respond to fluctuations in supply, thereby promoting investment in flexible technologies. If these projects can instead use EACs from a general carve-out to "top off" their operations and run around the clock, that flexibility incentive disappears, harming the industry's ability to adapt to a post-45V world when rapid responsiveness will be critical to capturing low-cost power in order to produce competitively-priced hydrogen.
- If Treasury nevertheless decides to include a general carve-out, it should be as small a percentage as possible, applied on per-generator and per-hour bases to limit the damage it causes, with the resultant energy attribute credits (EACs) preserving time-matching and deliverability requirements. It also should not be paired with more targeted exemptions, as the latter would eat into any logic (however flawed) of including the former.

- **To maintain emissions integrity, an exemption for capturing existing clean electricity that would otherwise be curtailed must include rigorous parameters,** such as exempting the incrementality component of EACs only during hours when the electrolyzers' nodal or zonal power price is below a pre-set threshold.
  - Treasury should still require electrolyzer operators to purchase EACs meeting the standard deliverability and time-matching requirements in these exempted hours.
  - Treasury should work with the U.S. Department of Energy (DOE) to determine the appropriate threshold locational marginal price (LMP)—such as $0 per megawatt-hour (MWh)—and consider setting different thresholds by region and periodically reevaluating these choices.
  - By default, an energy price-based exemption should only be available in organized power markets, but Treasury could work with independent third parties to potentially extend this exemption to other regions if they're equipped to validate and audit detailed grid modeling and hourly emissions forecasts from vertically integrated utilities.
- **Treasury should not offer an exemption intended to support clean energy resources that might retire if not for support from 45V.** There is little evidence this is needed to prevent retirements, and any attempt to capture such circumstances is far more likely to induce

substantial GHG emissions (by overestimating facilities' revenue requirements) and market manipulation. But if Treasury decides to create such an exemption, two options can help mitigate adverse impacts: an application-based approach and a formulaic approach.

- An application-based approach would require nuclear facilities to claim Civil Nuclear Credit (CNC) program funding from the Infrastructure Investment and Jobs Act (IIJA), the IRA's 45U tax credits that support existing nuclear generators, and any applicable state program funds to the fullest extent available, then demonstrate via a CNC-style application that additional support is needed to remain operational. Even then, the DOE should keep the qualifying generation as low as possible, though this allowance ought to remain for a full ten years.
- A formulaic approach would allow nuclear power plants to divert generation to electrolyzers any time their LMP falls below a facility-specific, DOE-determined threshold price that accounts for all other revenue sources (including from capacity markets, state subsidies, 45U, and the CNC program).

- **An exemption from incrementality for electricity in regions with binding emissions caps will maintain emissions integrity if and only if the caps are truly binding and resistant to shocks.** Ignoring the robustness and resilience of such a program could have the dual impact of being ineffective (i.e., worsening GHG emissions due to hydrogen production) and undermining its durability (e.g., driving policymakers to weaken or repeal the cap due to cost increases).
  - Specifically, Treasury should require emissions cap programs to have enforcement mechanisms that prevent 45-induced hydrogen production from leveraging allowance surpluses or triggering the release of extra allowances.
  - Treasury should also require emissions cap programs to control for leakage into unregulated regions or sectors to ensure GHG emissions induced by electrolyzers' operations are offset in full rather than merely in part (with such a determination likely requiring tighter deliverability regions and independent modeling).

ACKNOWLEDGEMENTS

The authors gratefully acknowledge Energy Innovation® colleagues Greg Alvarez, Anand Gopal, and Shannon Stirone for their review and communications contributions to this report. The authors would also like to acknowledge peer reviewers for their input, guidance, and contributions to this paper. The contents and conclusions, including any errors and omissions, are the sole responsibility of the authors, and peer reviewer input does not imply their support or endorsement of the work. Peer reviewers include Sara Gersen with Earthjustice, Hagan Han with Clean Air Task Force, and Jesse Jenkins with Princeton University.

Contents

INTRODUCTION

In its NPRM for the Section 45V Clean Hydrogen Production Tax Credit, Treasury proposes that electrolyzers must source power from new, local, time-matched clean power in order to earn the top tax credit tier of $3 per kilogram of hydrogen. Specifically, Treasury lays out three requirements:

- **Incrementality**, which requires using power from clean energy resources built no earlier than 36 months prior to the electrolyzer beginning commercial operations;
- **Deliverability**, which requires the electrolyzer and clean power to be sited within the same DOE National Transmission Needs Study region; and
- **Time-matching**, which requires the electrolyzer to draw clean energy to meet its operations on an hourly basis beginning in 2028 (with no grandfathering for projects built in prior years).

As discussed in Energy Innovation's April 2023 paper,[3] a recent DOE technical paper,[4] and a recent U.S. Environmental Protection Agency (EPA) letter,[5] these requirements are necessary to accurately verify the lifecycle GHG emissions intensities of electrolytic hydrogen needed to demonstrate compliance with the thresholds laid out in the IRA. They are also grounded in precedent and readily implementable. However, Treasury notes that "there are circumstances during which diversion of existing minimal (that is, zero or near-zero) emissions power generation to hydrogen production is unlikely to result in significant induced GHG emissions," with these falling outside of Treasury's core three-pillars framework.[6] Treasury describes three sets of such circumstances, involving the use of clean power that would come from:

- Generators that would otherwise **curtail** their output;
- Generators that would otherwise **retire**; or
- Generators located in states or regions that either have fully decarbonized power grids or **binding emissions caps.**

Due to the anticipated administrative challenges of designing rules tailored to each of these circumstances, Treasury asks whether it would be appropriate to "deem five percent of the hourly generation from minimal-emitting electricity generators (for example, wind, solar, nuclear, and hydropower facilities) placed in service before January 1, 2023, as satisfying the incrementality requirement," which we refer to as a "**general carve-out**" for existing clean energy.[7] Treasury ties this choice of 5 percent to the fact that the U.S. curtailed on the order of 5 percent of its renewable energy in 2022 and that about 5 percent of the existing nuclear power fleet may retire by 2032.

In this paper, we provide analysis aiming to support two categories of Treasury requests. First, we assess the merits of a general carve-out, addressing Treasury's question about whether the approach would appropriately balance "administrative feasibility and burden with accuracy of identifying circumstances with a low risk of induced grid emissions."[8] Second, we provide implementation options for targeted approaches that would cover (1) clean power that would otherwise be curtailed, (2) avoided facility retirements, and (3) states or regions with binding emissions caps; we also note how even these options carry some risk of adverse outcomes yet, in some cases, may strike the right balance between accuracy and administrability.

"GENERAL CARVE-OUT" APPROACH FOR EXEMPTING EXISTING CLEAN ENERGY

Treasury is rightly seeking options to qualify existing clean energy resources in any circumstances in which their use wouldn't cause an increase in GHG emissions elsewhere on the grid. It also is appropriately asking how to do so with as few administrative constraints as possible. However, a general carve-out for five—or any—percent of existing clean energy to qualify without meeting specialized criteria would unfortunately miss Treasury's target by a wide margin, failing to even come close to its admirable intent. This section explains why a general carve-out would be a sure-fire mismatch for capturing power that would otherwise be curtailed, retire, or not induce grid emissions due to a binding emissions cap. While we find that a general carve-out of any kind would materially undermine 45V's emissions requirements, we highlight design choices within this unworkable framework that would exacerbate or reduce relative impacts.

A GENERAL CARVE-OUT WOULD FALL FAR SHORT OF ITS INTENT

Curtailment

Curtailment generally refers to the waste of renewable energy due to needing to always keep electricity supply and demand in perfect balance. When considering a given power grid or region, curtailment can be split into two parts: (1) system curtailment, in which there is too much available renewable electricity relative to a regional grid's total system demand in a given interval (i.e., the next increment of demand anywhere on the regional grid could be served by existing renewable generation that is being curtailed); and (2) local curtailment, in which localized transmission congestion prevents excess renewable electricity from being absorbed by the system at a specific place on the grid even if a customer somewhere else might be interested in using it.

While we would ideally waste as little renewable electricity as possible, curtailment does serve a purpose in running economically efficient power markets. It sometimes makes more sense to waste power in order to optimize supply and demand for lowest cost across a sprawling, complex electricity system.[9] However, by strategically siting new electricity demands like batteries and electrolyzers, it's possible to make use of at least some of the renewable electricity that would otherwise go to waste.

Electrolyzers at today's prices do not have a business case running on curtailed power alone. Electrolyzers have two main costs: the capital investment associated with equipment, and the operational costs, of which the largest component is electricity. Until the capital costs fall (which many expect will occur as the industry scales), electrolyzers will have to run most of the time in order to minimize their per-unit capital costs and maximize revenue from the 45V tax credit.

However, by its nature, curtailment is too **spiky in time** and **dispersed in location** for electrolyzers to ever directly capture more than a small fraction of this would-be wasted power. Therefore, under a general carve-out exemption that allows for the indiscriminate consumption of existing clean energy, it's almost guaranteed that electrolyzers would mostly draw power in hours when curtailment is very low or zero. This means they would primarily induce fossil fuel power plants to provide more power to the grid rather than meaningfully soak up excess renewable electricity. In other words, a general carve-out would subsidize hydrogen production that is, on average, highly emissions intensive.[iv]

California example

Consider the California Independent System Operator (CAISO), an organized power market where a high penetration of wind and solar energy resources makes it a hotbed for curtailment. A plot of CAISO's 2023 total curtailment (i.e., system plus all local curtailment) reveals the scattershot nature of what Treasury is hoping to capture under a general carve-out.

As shown in Figure 1, nearly a quarter of the total wind and solar energy wasted in CAISO in 2023 occurred in just 1 percent of hours, and nearly 80 percent of that waste happened in 10 percent of hours. 100 megawatts (MW) of electrolyzers running only when excess renewable energy was available—assuming these electrolyzers magically had access to all CAISO power that would have otherwise been curtailed regardless of congestion in the transmission system—would only manage a 34 percent load factor.[v] Even this meager runtime plummets as you increase scale—an electrolyzer buildout of 1,000 MW drops their collective load factor to a mere 16 percent. Despite having access to free electricity supported further by lucrative 45V tax credits, this load factor wouldn't pencil out with today's expensive electrolyzers.

[iv] While the following examples use data from recent years, the conclusions of this analysis would remain true even in a world with far more curtailment—that is, the nature of curtailment as spiky in time and dispersed in location keeps the vast majority of it from being accessible to technologies that need frequent operations to be financially viable.

[v] An electrolyzer's "load factor" is the amount of hydrogen it produces in a year divided by the amount of hydrogen it could theoretically produce in a year if it ran at full, continuous operations.

Figure 1. CAISO daily total curtailment (2023)

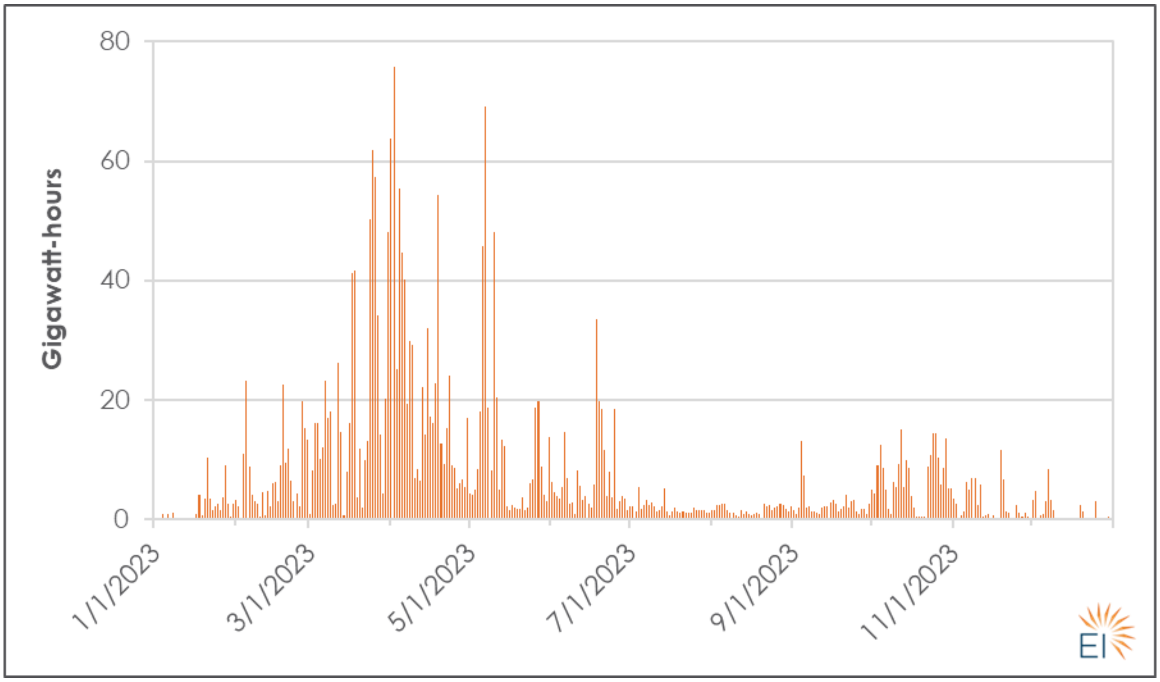

Most of CAISO's wasted wind and solar energy in 2023 occurred in a small subset of days, making it difficult to capture.

Unfortunately, the on-the-ground reality is even worse. While electrolyzers anywhere in CAISO can access power that would be curtailed at the system level, capturing it at the local level requires being sited within the specific part of the network that is facing export congestion (i.e., with transmission lines overloaded, renewable generators have nowhere to send their excess power)—and CAISO doesn't provide more granular data on where local curtailment is happening.[10] However, it is certainly not all happening in the same place.

This means that while any electrolyzer might be located such that it can access more than just system curtailed power by tapping into locally curtailed power, no single unit can be everywhere at once, able to access total curtailment (the combination of system plus all aggregated local curtailment). In other words, while system curtailment represents the minimum amount of wasted power that any electrolyzer can access (a lower bound), we don't know what the upper bound consumption of curtailed power for an electrolyzer at a specific location is—but even for the best chosen locations, it will be far below total curtailment.

A plot of CAISO's 2023 system curtailment reveals that it not only makes up a small share (about a fifth) of total curtailment—it also comprises mostly single-hour spikes. More than 70 percent of system curtailment occurred in 1 percent of hours, and more than 99 percent happened in 10 percent of hours. A mere 100 MW of electrolyzers would top out at a 7 percent load factor if relying only on system curtailment—clearly unworkable for financing new projects, even with far lower capital costs.

Figure 2. CAISO daily system curtailment (2023)

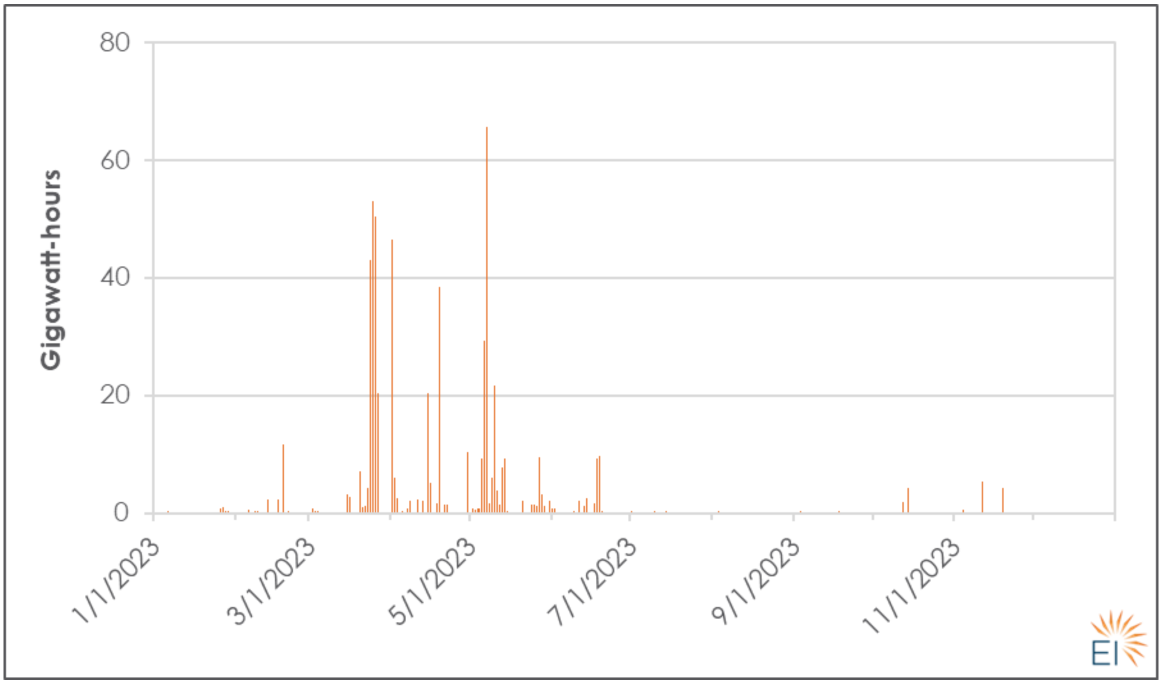

The CAISO system was heavily oversupplied (thereby curtailing wind and solar energy) in just a few days in 2023.

Let's now consider how a general carve-out might play out under the assumption that it would help unlock access to curtailed power. A general carve-out would allow existing clean energy resources to qualify 5 percent of their generation for EACs that need not meet the incrementality criterion (but must still comply with deliverability and time-matching). If an electrolyzer purchases these EACs in hours when (or locations where) its operations are not using clean energy that would otherwise have been curtailed, it is very likely to induce GHG emissions elsewhere on the grid. This emissions impact is worsened the wider this discrepancy gets.

To investigate the size and impact of this discrepancy under a general carve-out, we compared the hours in which CAISO experienced actual curtailment in 2023 with the availability of EACs that would qualify under a general carve-out. The figure below ranks system and total curtailment by hour, acting as lower and extreme upper bounds for what curtailed power would be available to electrolyzers in CAISO in 2023. It also shows the amount of existing clean energy that would qualify under a 5 percent general carve-out applied on a per-hour basis (without the ability to reshuffle this allowance into other hours).

Figure 3. CAISO ranked curtailment vs. 5 percent general carve-out (2023)

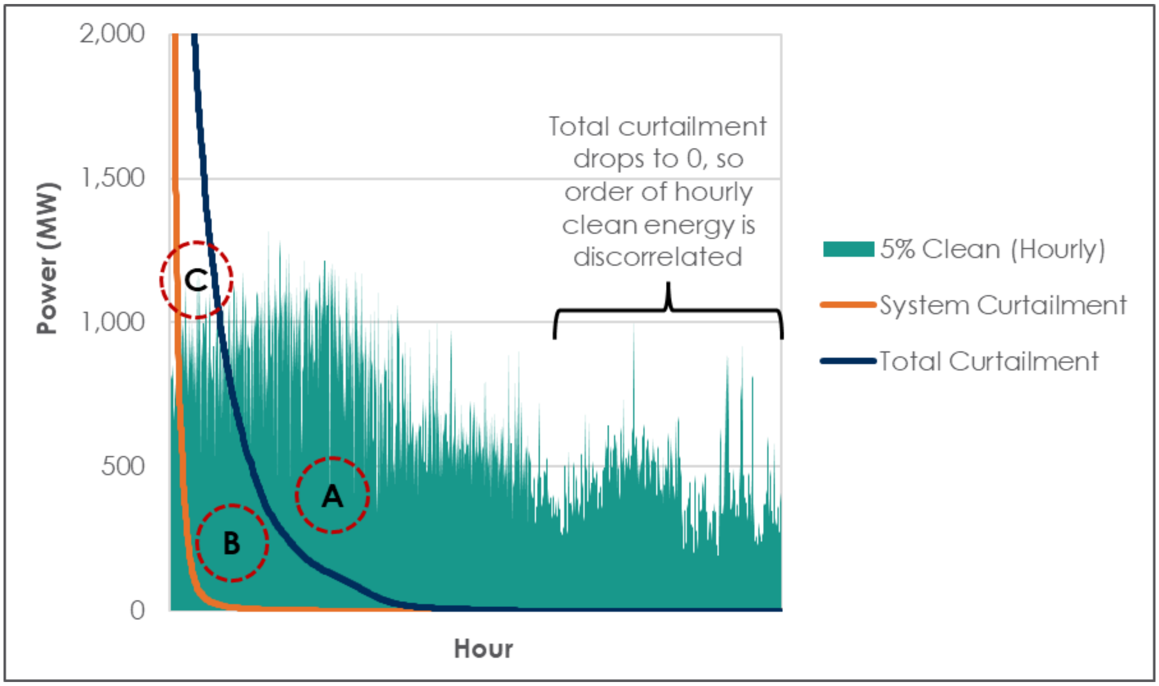

The data used for the "5% Clean (Hourly)" series align with the same hours as the ranked "Total Curtailment" series but becomes discorrelated (effectively random) when "Total Curtailment" drops to zero. The "System Curtailment" series is independently ranked. The y-axis cuts off to make the rest of the figure legible, but system and total curtailment reach on the order of 9,000 MW in some hours.

The figure above helps us visualize three areas of interest loosely marked by red dashed circles: (A) the area above the curtailment curves but below the 5 percent clean indicator, showing the huge volume of EACs that a general carve-out would mistakenly qualify for 45V (false positives); (B) the area below the curtailment curves and below the 5 percent clean indicator, showing the comparatively much smaller volume of EACs that a general carve-out would incidentally accurately qualify for 45V (true positives); and (C) the area below the curtailment curves but above the 5 percent clean indicator (mostly not visible), representing the significant volume of curtailed power that, if captured, would make sense to qualify for 45V but would not be eligible under a general carve-out (false negatives).

We also consider the impact of allocating the general carve-out on an annual basis. To simplify the analysis, we assume the presence of electrolyzers that would run strictly on exempted EACs, which would purchase them in equal quantities every hour to run continuously and maximize 45V credits.

The reality may be much worse for emissions. For example, consider a three-pillars-compliant project that draws most of its energy from new wind and solar resources. In times of significant curtailment, the electrolyzer is likely already operating at max capacity due to high output from its dedicated renewables, meaning it can't ramp up further to take advantage of this wasted power with exempted EACs from a general carve-out. In times of zero curtailment, the electrolyzer may not have access to enough incremental generation to run at full capacity and therefore would seek EACs (paired with purchases of grid power) to boost its load factor. Such an outcome would drive a higher GHG emissions intensity (i.e., a greater proportion of EACs in area A than area B in Figure 3) than our assumption of equal EAC availability in all hours.

Table 1 summarizes the results for a 5 percent general carve-out, illustrating differences between having access to total vs. system curtailment and applying the carve-out on an annual average vs. hourly basis. Table 2 summarizes the results for the same analysis but with a 10 percent general carve-out.

Table 1. Key indicators from 5 percent general carve-out (CAISO 2023)

Scenario	[A] Share of exempt EACs that are dirty	[B] Share of exempt EACs that are clean	[C] Share of curtailed power that is not captured	Average EAC emissions intensity ($kgCO_2/kgH_2$)
Total curtailment Annual average carve-out	80%	20%	58%	18.0
Total curtailment Hourly carve-out	77%	23%	50%	17.2
System curtailment Annual average carve-out	97%	3%	67%	21.7
System curtailment Hourly carve-out	96%	4%	63%	21.6

Table 2. Key indicators from 10 percent general carve-out (CAISO 2023)

Scenario	[A] Share of exempt EACs that are dirty	[B] Share of exempt EACs that are clean	[C] Share of curtailed power that is not captured	Average EAC emissions intensity ($kgCO_2/kgH_2$)
Total curtailment Annual average carve-out	86%	14%	39%	19.3
Total curtailment Hourly carve-out	84%	16%	32%	18.9
System curtailment Annual average carve-out	98%	2%	53%	21.9
System curtailment Hourly carve-out	97%	3%	48%	21.8

We can draw a few conclusions from this analysis:

- A general carve-out may allow electrolyzers to incidentally capture some share of clean energy that would otherwise be curtailed, but the high false positive rate for such a design would result in substantial induced GHG emissions—a real-world average annual rate of hydrogen production emissions of on the order of 30 to 45 times above 45V's emissions intensity threshold for earning the $3/kgH$_2$ credit and 1.5 to two times worse than hydrogen produced via unmitigated steam methane reformation. (This impact could be even more severe in regions where coal power plants are on the margin.)
- Distributing this carve-out on an hourly rather than annual average basis helps ensure more of its allowance aligns with clean energy that would otherwise be curtailed, but it does not do so nearly enough to justify the overall concept.
- Increasing the general carve-out allowance (e.g., from 5 to 10 percent) may capture a greater magnitude of clean energy that would otherwise be curtailed, but it opens the door to a much larger proportional share of fossil fuel energy consumption and drives higher average annual GHG emissions rates of hydrogen production emissions.

Texas example

As one more example, let's consider the ERCOT power market in Texas, which (unlike California) has much more wind than solar—a better match for maximizing electrolyzer load factors. In lieu of granular curtailment data, a comparison of relative zonal power prices reveals the risk of employing a general carve-out approach for attempting to use clean power that would otherwise be curtailed.

Power markets function by arranging the cheapest-available generators to meet electricity demand in any given moment, with power plants bidding their marginal costs to a central operator. Renewable energy facilities like wind and solar have negligible marginal costs, making it rational to bid $0/MWh into the market. Renewables will even bid below zero if they earn production tax credits for each MWh that they generate. Facilities that burn fuel, like coal, natural gas, and nuclear, generally have higher (positive) marginal costs to account for the price of fuel as well as the variable operations and maintenance needed to manage the plants' performance. Thus, negative wholesale prices are a useful and relatively reliable signal for periods when clean energy power plants are likely curtailing some of their output.

Figure 4 plots all 15-minute real-time energy prices in ERCOT in 2023 for the West Zone and Houston Zone. The data reveal that, in 2023, the West Zone experienced negative power prices in 8 percent of its intervals while the Houston Zone experienced such prices in just 2 percent of intervals. This gap suggests a much higher incidence of curtailment in the West Zone. However, the Houston Zone had a slightly lower average power price below $80/MWh (or the price below which a typical electrolyzer would choose to make hydrogen if it had access to the 45V tax credit and a sale price of $1/kg) than the West Zone.[vi] Thus, with access to a general carve-out, developers would be incentivized to build electrolyzers in Houston—especially given the city's proximity to traditional hydrogen consumers that would reduce hydrogen transportation costs.

Figure 4. ERCOT West Zone vs. Houston Zone 15-minute real-time energy prices (2023)

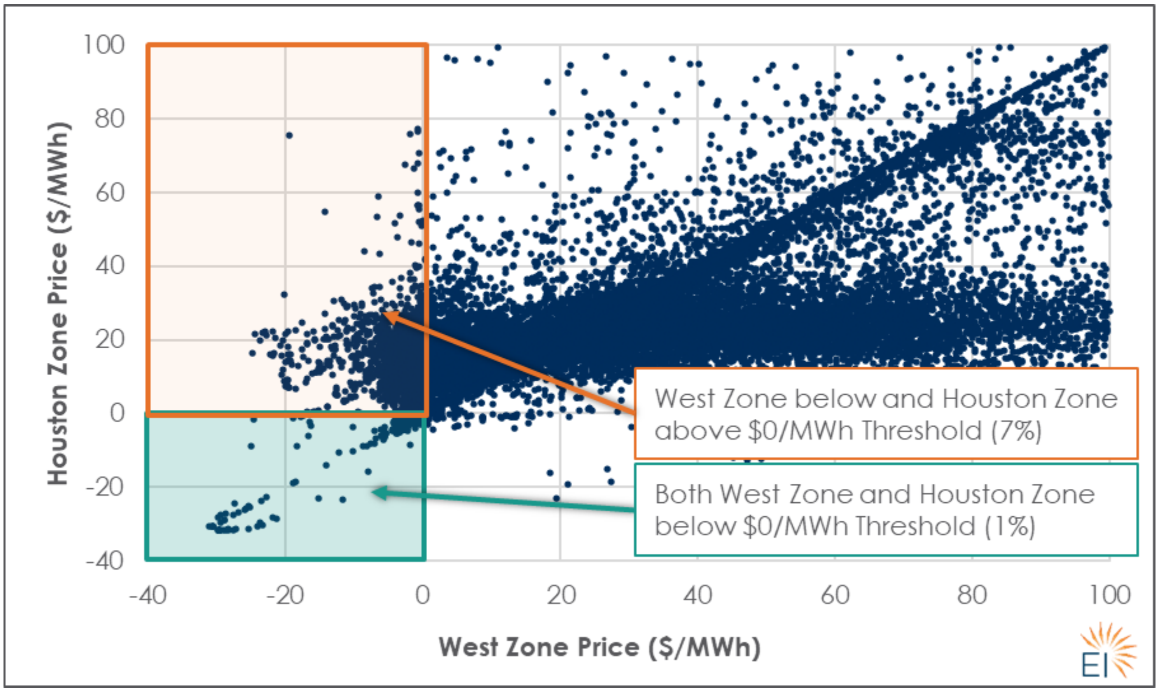

Electrolyzers built in the Houston Zone would not be able to access most of the curtailment occurring in the West Zone due to transmission congestion.

This mismatch in the locations of electrolyzers and curtailed power within the broader ERCOT market creates a problem similar to what we demonstrated for California. Our proxy for curtailment, a negative LMP in the West Zone,[vii] creates a lot of false positives for a general carve-out exemption: intervals where it seems like excess clean power in the West Zone might be available to Houston Zone electrolyzers but is not actually deliverable due to transmission congestion (as indicated by coincident positive LMPs in Houston).

[vi] Specifically, Houston Zone and West Zone average power prices in intervals below $80/MWh were $23.63/MWh and $24.79/MWh in 2023, respectively. Average power prices over all hours were $49.75/MWh for the Houston Zone and $59.42/MWh for the West Zone, though electrolyzers would not run beyond a certain breakeven price.
[vii] There were only 89 intervals in which the Houston Zone experienced negative LMPs while the West Zone experienced positive LMPs in 2023—so the vast majority of available curtailment in this example is indicated by West Zone LMPs.

Specifically, Figure 4 reveals how Houston-area electrolyzers would capture far less curtailed power in reality due to interzonal transmission congestion. That is, the West Zone and Houston Zone only experience negative prices together in 1 percent of all 15-minute intervals (true positives); in another 7 percent of intervals, the West Zone experiences negative prices, but the Houston Zone's higher prices indicate the presence of interzonal congestion (false positives). In these intervals, the general carve-out would allow the Houston-area electrolyzer to earn tax credits while being unable to access West Zone curtailment, instead likely inducing a local fossil fuel power plant to ramp up to serve its demand. Even if pricing patterns change over the coming decade, this example is indicative of a general carve-out's failure to incentivize the behavior Treasury intended with its conception of this idea.

Retirement preventions

A general carve-out would do a similarly poor job of aligning with existing nuclear power plants that would otherwise retire if not for financial support via 45V. This thought experiment is much more straightforward. All nuclear facilities can essentially be divided into two categories: those that will retire without a new source of revenue (such as 45V), and those that are doing fine but would take 45V—in at least some if not all hours—if offered.

The average revenue needed by an existing nuclear power plant is on the order of $45/MWh,[viii] while hourly wholesale power prices fall well below $80/MWh in most hours. This means that while the group threatened by retirement would gain more from earning 45V credits, the financially viable group would still pursue 45V to the fullest extent possible given that it represents an essentially risk-free source of revenue higher than what it can typically earn from power markets.

The general carve-out is predicated on supporting the roughly 5 percent of existing nuclear facilities that may retire without extra financial support by 2032, but it doesn't discern between these two groups. So the vast majority of this exemption would flow to plants that do not need it—thereby raising GHG emissions substantially—while the plants that might need it to stay in the market might get too low of an allowance from the carve-out to even be effective at keeping the facility viable as intended.

Binding emissions caps

Lastly, a general carve-out would be a poor approach for supporting existing renewable energy facilities in states with binding emissions caps, as any use of the carve-out for this purpose outside

[viii] We base this on costs reported by the Nuclear Energy Institute (https://www.nei.org/CorporateSite/media/filefolder/resources/reports-and-briefs/2023-Costs-in-Context.pdf) and The Brattle Group (https://www.brattle.com/wp-content/uploads/2017/10/5727_brattle_nuclear-carbon_whitepaper_-_dec2016.pdf) as well as Congress's decision to cap the 45U production tax credit for existing nuclear power plants at total revenues of $45/MWh.

of such states would not be a match. In 2023, 14 states[ix] had CO_2 emissions trading schemes, meaning most of the U.S. could still access a general carve-out despite having no cap.[11] Yet the reality is much worse—as we discuss later in this paper, even these 14 states might not be able to call their caps "binding," meaning Treasury cannot assume electrolyzer GHG emissions would be offset elsewhere without further investigation or guardrails.

A GENERAL CARVE-OUT WOULD HARM CLIMATE AND INNOVATION

The consequences of a general carve-out would be severe. Exempting 5 percent of existing clean energy from the incrementality requirement would allow for on the order of 1.5 MMT of dirty electrolytic hydrogen production per year, contributing roughly 30 to 60 MMT of CO_2 emissions annually. Figure 3 shows the region-level hydrogen production emissions intensity from ignoring an incrementality requirement—including from a mismatched exemption like a general carve-out—using the EPA's AVERT tool for assessing marginal GHG emissions rates.

Figure 5. Short-run marginal GHG emissions impacts of forgoing additionality (2021)[12]

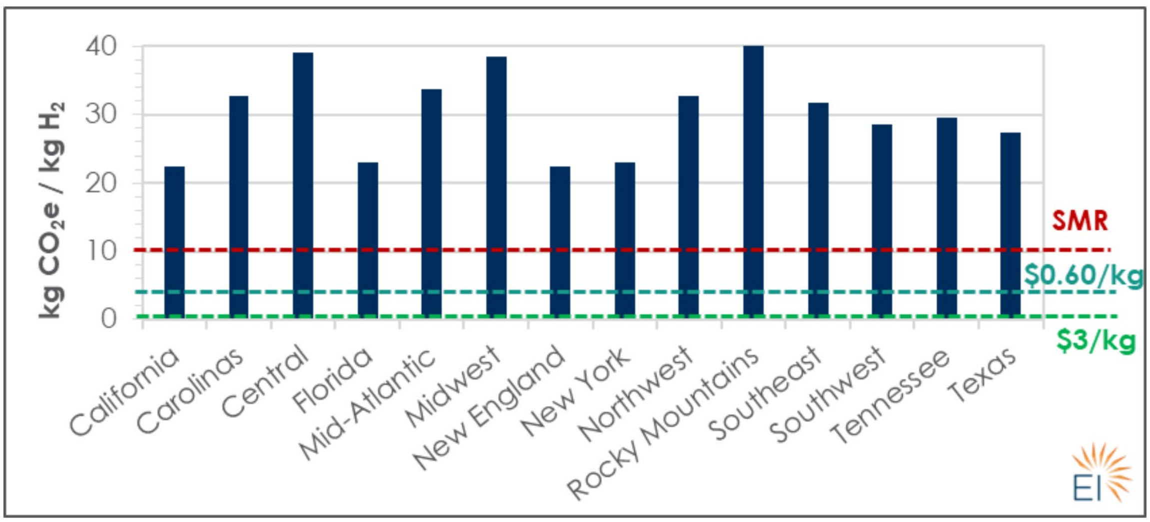

Without additionality, electrolyzers everywhere in the U.S. would cause GHG emissions that are at least twice as high as those from steam methane reformation. Analysis uses AVERT with 2021 data. It is not representative of emissions impacts in later years and does not assess long-run marginal GHG emissions impacts.

Modeling by Rhodium Group arrives at a similarly alarming outcome: that a 5 percent general carve-out could "cause a huge increase in net system-wide emissions—up to nearly 1.5 billion

---

[ix] These states were California, Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Oregon, Rhode Island, Vermont, Virginia, and Washington.

metric tons of increased emissions cumulatively through 2035" if electrolyzer operations induced more generation from the dirtiest power plants on the grid.[13]

While these emissions outcomes are harmful in themselves (and contravene requirements of the statute), a general carve-out could cause much more damage to the extent it hampers the incentive to invest in innovation. Consider an electrolyzer project that predominantly sources its electricity from incremental clean energy resources. A general carve-out could allow such a project to "top off" its load factor with electricity from existing clean energy resources. This could either bump it up to running around the clock (removing the need to invest in hydrogen storage) or at least make its ramping needs much more predictable or easy to manage (dampening the incentive to invest in flexible electrolyzers).

As discussed in our earlier paper on 45V, an industry built on inflexible electrolyzers, a dearth of hydrogen storage, and rigid business models will be set up to fail or require perpetual subsidy extensions.[14] Although a general carve-out would only support a relatively small amount of hydrogen production, this could be detrimental to the industry's long-term viability to the extent this production gets spread across many more facilities.

LIMITING THE DAMAGE FROM A GENERAL CARVE-OUT

In any conceivable scenario, the harm of a general carve-out far outweighs the good, and we recommend Treasury disregard this option in favor of targeted approaches. However, if Treasury does offer an indiscriminate, broadly applied exemption, several design choices can mitigate the damage.

- Keep the percentage of existing clean energy that qualifies for this exemption as low as possible.
- Do not offer both a general carve-out and targeted approaches for capturing curtailment, retirement preventions, and binding emissions caps, as the latter option set would further reduce any of the slight value the former might incidentally provide.
- Maintain deliverability and time-matching requirements for EACs that are exempted from incrementality under a general carve-out.
- Apply the general carve-out on a per-generator basis without permitting the trade of exempt EACs between facilities, as this would at least ensure some share of the carve-out continues to support flexible electrolyzer operations (e.g., by matching variable output from wind and solar resources).
- Apply the general carve-out on a per-hour basis (rather than per day, month, or year), as this would similarly help train electrolyzers for flexible operations and encourage more investment in hydrogen storage to smooth gaps in production. This would also prevent generators from shifting exempt EACs to hours when they are most valuable (i.e., scarce),

which would have the adverse impacts of aligning electrolyzer operations with hours when fossil fuels are most likely to ramp up to meet demand and undercutting the incentive to build new clean energy to serve these needs.

TARGETED APPROACHES FOR EXEMPTING EXISTING CLEAN ENERGY

Because a general carve-out approach is ill-equipped to estimate GHG emissions from electrolytic hydrogen production with any reasonable degree of accuracy, Treasury is also requesting comments on more targeted approaches that can qualify electricity from existing clean energy resources under appropriate circumstances without facing significant administrative complexity. This section offers ideas for how Treasury can design exemptions that do a far better job at aligning the use of electricity from existing clean energy resources with conditions that would prevent inducing GHG emissions elsewhere on the grid.

CURTAILMENT

In our original paper on 45V, we state that "an additionality framework should ideally have a mechanism to give credit to capturing otherwise-curtailed clean electricity."[15] Below, we walk through the highest-fidelity, administrable approach for implementing such an exemption. We also discuss the adverse consequences that could still result from such an approach, concluding that if proper safeguards are put in place, the risks of potential emissions increases would be adequately balanced with the rewards of qualifying more electricity that would not induce upstream emissions.

Background

In modern U.S. power markets, generators receive a price for their power according to their local marginal price (LMP) node. The LMP represents the value of electricity at that node, including all the system costs to adjust other generators' operations in order to deliver that power at the lowest-possible cost while keeping the grid stable.

Most markets aggregate prices from LMP nodes (which figure in the hundreds) to a handful of zones or hubs. In general, zones are mostly for buyers, and hubs represent some kind of averaged price for traders.

When negotiating a power purchase agreement, a key issue for the participants in the contract is the "basis risk": generally, the difference in geographic price (basis) between the LMP and the hub or zone. Traders don't want to track hundreds of nodes—they want to trade at the hub or zone—so the generator typically takes on the basis risk. This usually means they accept the risk that they will get paid less at their LMP node than they will receive from a trader or customer at the hub or zone. If this spread is large, that is a sign of congestion on transmission paths from the generator to a generic class of buyers in the larger hub or zone that the generator's node exists within.

Very low LMPs, such as below zero, generally indicate that clean energy resources are curtailing some of their power. So in the context of 45V, such LMPs can be a decent proxy for low to zero marginal emissions caused by an electrolyzer's coincident operations—though we'll discuss later why this is not always the case.

Whether power would be curtailed depends on the behavior of both buyers and sellers. Each can react to prices in a way that affects demand. Next, we examine whether the buyer's or seller's perspective is more useful in this counterfactual based on prices, concluding that prices facing the electrolyzer are a better proxy for power curtailment.

Generator point of view on curtailed power

First, let's explore the point of view of an existing clean energy resource. Ideally, if its power is being curtailed during some interval, it would make sense to make it eligible for the 45V credit for an electrolyzer that can actually make use of it. Negative LMP prices are a system signal that even zero marginal cost generators should dial down their power output.

The problem with qualifying existing clean energy in this manner is that it may not be deliverable to an electrolyzer in a generic location inside that zone or region due to transmission constraints. LMPs can go negative even when the zonal price is higher, specifically due to intrazonal transmission congestion—the power available from clean energy generators exceeds what can flow off the local node. If such a difference exists, running an electrolyzer may induce fossil fuel power plants to ramp up elsewhere on the system rather than soak up otherwise curtailed clean energy. Thus, Treasury should not look to a generator's circumstances as a means for identifying existing clean energy that ought to receive an incrementality exemption.

Electrolyzer point of view on curtailed power

An alternative exemption pathway would be to use an LMP criterion tied to the actual nodal location of the electrolyzer hoping to receive the 45V tax credit. This has the advantage of being more consequential: the local LMP really does reflect the estimated real-time price impact (and hopefully the emissions impact by proxy) of the electrolyzer's operations. Increased use of the electrolyzer will raise the LMP, ideally signifying that it is making use of clean power in this locality that would otherwise have been curtailed. An electrolyzer meeting this LMP threshold—for example, of buying negatively priced power—would be permitted to disregard the incrementality component of EACs for the energy consumed in that interval.

One obstacle to this buyer-side LMP approach is that electricity buyers typically are not buying power at a node and may not even be able to do so under existing tariff designs. However, this approach can still work in two ways.

First, electrolyzers can opt to use the zonal price that they're actually paying for power. Zonal prices won't fall below zero as often as specific LMPs (since they would only do so if the aggregate of nodal prices were sufficiently low), so electrolyzers would have far fewer opportunities to capitalize on curtailed power, but it would be administratively simple.

Second, the market operator could assign electrolyzers to their physically relevant LMP node, as an electrolyzer's consumption will show up at that node in the market operator's model. In intervals when they fall below the LMP threshold, any power that electrolyzers buy should induce low or zero upstream emissions, so there is no need to verify a specific commercial transaction at the node.

Risks and mitigation measures

A buyer-side approach to qualifying clean energy that would otherwise be curtailed is appealing, but it comes with some risks. We discuss three of these risks in turn below, along with actions Treasury can take to mitigate them where applicable.

The first risk is the **double-counting** of clean energy. If an electrolyzer actually reduces curtailment, as is the intended outcome, then a clean energy resource somewhere is generating more power and therefore more EACs. If the electrolyzer is allowed to qualify for 45V credits under low LMP conditions without buying EACs, then that generator will sell them to another counterparty—say, to an electrolyzer elsewhere in the deliverability region (if the clean energy resource is "incremental") or to a utility or corporation that can apply it as a standard, annual renewable energy credit against the actual use of fossil power (if the clean energy resource is of an older vintage, i.e., "existing"). This is double-counting: the electrolyzer is implicitly claiming the use of clean energy by virtue of qualifying for 45V, while the purchaser of the EAC claims the use of clean power that it's not actually consuming.

Fortunately, this risk has a straightforward solution. If Treasury adopts this approach, it should still require the electrolyzer to buy deliverable, time-matched EACs, even though the electrolyzer need not meet the incrementality criterion.

The second risk is of **low LMPs not accurately reflecting local marginal emissions (LMEs)**. For example, in a market like Southwest Power Pool with vertically integrated utilities and a higher share of inflexible coal power plants, self-scheduling is a common issue.[16] In short, utilities can decide to run fossil power plants at a loss, including down to or below $0/MWh, in order to either avoid frequent ramping or create conditions more beneficial to the utility's overall fleet (e.g., unlocking more hours of conditions that qualify for this exemption for electrolyzers that it also owns). It's therefore possible to see conditions in which electrolyzers buy power at negative LMPs with coal plants providing this power.

This risk can be lessened by setting a low LMP threshold for determining when this exemption applies. The tradeoff is tricky: a higher threshold helps electrolyzers capture more clean energy that would otherwise be curtailed or have no induced GHG emissions impact. But it also opens the door to more instances of fossil fuel generation sneaking in. The gaming risk is especially concerning: $3/kgH$_2$ is worth on the order of $60/MWh, so coal-fired power plants could take quite a financial hit (i.e., by running more often to lower their LMPs below the threshold) if their owners also had sufficient electrolyzers nearby to benefit from more tax credits. To find the right balance, Treasury should work with the DOE to choose a threshold LMP and also should consider setting different thresholds by region and periodically reevaluating these choices.

This risk can also be mitigated to a greater degree (albeit with a potentially higher administrative burden) by working with independent experts capable of setting guardrails on induced emissions impacts, such as through using an independent system operator's (ISO) or regional transmission organization's (RTO) LME data.[x] Such third parties can opt in to provide this higher-fidelity information. Partnerships with third parties can also open the door to exempting curtailment in regions without organized power markets, particularly if they're equipped to validate and audit detailed grid modeling and hourly emissions forecasts from vertically integrated utilities. Greater emphasis in this rulemaking on hourly time-matching will likely produce a side benefit of driving further innovations in marginal emissions accounting that will accommodate this practice.

The third risk is purposefully taking advantage of the conditions that cause curtailment in order to arbitrage the difference in value between EACs of older and newer vintages, which we deem **vintage-washing**. In this case, a developer could build new renewables and electrolyzers in a pocket (or area of the transmission grid with limited export capacity) with lots of existing congestion. In periods of low LMPs, the exempted electrolyzer could buy lower-value, older-vintage EACs from other existing clean energy resources outside the congested pocket, while the new clean energy resources could sell their higher-value, newer-vintage EACs to other buyers in the much larger DOE National Transmission Needs Study region. It's not clear to what degree developers would find such a scheme attractive, nor is there a clear way to prevent this, but it reveals the type of gaming that could arise with even a carefully designed exemption.

Finding the right balance

In seeking an approach to qualify the use of existing clean power that would otherwise be curtailed, Treasury must balance practical concerns and administrability with accurately rewarding hydrogen production that meets 45V's stated emissions thresholds. However, there will likely be few instances in which electrolyzers supported by 45V would legitimately capture clean power that otherwise would have been curtailed—at least relative to the volume of electricity they will consume from incremental clean energy resources.

For example, as discussed earlier in greater detail, electrolyzers built during the lifetime of the 45V tax credit are unlikely to be financed strictly through relying on a curtailed power exemption. Today's electrolyzer capital costs (particularly for flexible technologies capable of reacting quickly to operate only in these intervals) and lower curtailment volumes suggest electrolyzers will need to primarily rely on incremental clean energy resources to achieve a high-enough load factor to

---

[x] Resurity offers an LME product (https://resurety.com/resurety-and-watttime-partner-to-increase-accessibility-to-high-quality-marginal-carbon-emissions-data/), and the largest organized power market, PJM, also offers some LME data (https://www.pjm.com/-/media/etools/data-miner-2/marginal-emissions-primer.ashx).

reach financial viability. Any otherwise-curtailed power would likely only provide relatively small and less predictable increments of clean power to electrolyzer projects.

Further, projects supported by incremental wind and solar resources will often already be operating at full capacity in the same hours that existing clean energy resources are curtailing some power. This is because renewables in a local geographic area experience solar incidence or wind speeds simultaneously, and the incremental wind and solar resources would likely be sized to allow the electrolyzer to produce hydrogen at full capacity without needing to rely on clean power that would otherwise be curtailed.

Thus, while it makes sense to pursue an option to exempt existing clean energy that would otherwise be curtailed, Treasury must take great care to not tip the scales too far toward ease of access and away from integrity of verifying GHG emissions impacts. Doing so could open the door to far worse repercussions, including through standard market operations (e.g., inflexible fossil fuel power plants running at a loss during low LMP hours to avoid the costs of ramping) and intentional gaming (e.g., market participants strategically self-scheduling fossil fuel facilities' operations to unlock more exempted electricity for their electrolyzers). In other words, without sufficiently tight guardrails, this exemption runs the risk of inducing far more GHG emissions than it rewards through the actual capture of existing clean energy that would otherwise be curtailed.

Setting reasonable restrictions for this exemption may leave some curtailment on the table; indeed, this is basically guaranteed, as even the loosest rule set would still not lead electrolyzers to capture all such would-be wasted clean power. Fortunately, Treasury's strong three-pillar framework sets up the conditions in which post-45V electrolyzers can benefit from a larger pool of excess clean energy.

Specifically, electrolyzers that have rolled off the 10-year tax credit—as well as electrolyzers built after 45V expires—will have far lower capital costs than today's electrolyzers. They will also exist in a world with far greater clean energy penetration and, by extension, curtailment. These subsidy-free electrolyzers will need to situationally buy very cheap power (below $20/MWh) to make competitively priced hydrogen. Thus, strong rules help build the type of technology base that will capture higher shares of curtailment down the road and keep the industry financially viable in a post-subsidy world.

Takeaways and recommendations

Our assessment of potential exemptions for curtailed power can be summarized as follows:

- Treasury should only offer an exemption for capturing existing clean electricity that would otherwise be curtailed if it sets rigorous parameters, such as exempting the incrementality component of EACs only during hours when the electrolyzers' nodal or zonal power price is below a pre-determined threshold.
    - Treasury should still require electrolyzers to purchase EACs that meet the standard deliverability and time-matching requirements in these exempted hours.

- Treasury should work with the DOE to determine the appropriate threshold LMP (such as $0/MWh) and should consider setting different thresholds by region and periodically reevaluating these choices.
- Treasury should seek to employ independent experts such as ISO/RTOs to further guard against adverse outcomes (e.g., self-scheduling of fossil fuel power plants during low LMP hours).
- By default, this exemption should only be available in organized power markets, but Treasury could work with independent third parties to potentially extend this exemption to other regions if they're equipped to validate and audit detailed grid modeling and hourly emissions forecasts from vertically integrated utilities.

- This proposal represents a reasonable balance between allowing a significant share of otherwise-curtailed clean power to qualify and using readily administrable guardrails to protect against the bulk of this approach's unintended consequences and gaming risks.
- Regardless of the rule set, there are relatively few instances in which 45V-subsidized electrolyzers would legitimately use clean energy that would otherwise be curtailed. Loosening rules for this exemption is much more likely to increase the share of fossil fuel power being used by electrolyzers relative to the share of otherwise-wasted clean power. By contrast, a strong framework tightens the scope of the exemption to what ought to qualify; it also helps train electrolyzers to capture clean, cheap power long after 45V expires, setting the industry up for sustainable growth in a post-subsidy world.

RETIREMENT PREVENTIONS

In our original paper on 45V, we state that an exemption targeting clean energy facilities that would otherwise retire is likely unnecessary since such an exemption would only reasonably apply to nuclear power plants, yet these facilities are already well supported by other federal and state policies.[17] Our conclusion remains the same. However, should Treasury choose to provide an exemption targeting facilities facing real or imagined retirement risk, we offer two frameworks intended to limit adverse emissions-leakage outcomes: an application-based approach and a formulaic approach.

Application-based approach

The application-based approach would require clean energy facilities to meet a standard similar to the IIJA's CNC program—an already-established financial test that can provide the needed evidence to reasonably judge whether a facility is truly at risk and deserves an exemption for some share of its output.

Under such an approach, Treasury should require that (1) the CNC program's funding be exhausted (since the program is intended to keep nuclear power plants financially viable); (2) nuclear facilities use to the fullest extent the IRA's 45U tax credit (which provides up to $15/MWh in tax credits); and (3) nuclear facilities exhaust any funding from relevant state policies. Should a nuclear power plant be able to prove that it still needs additional support to remain operational, Treasury could consider extending an exemption.

As for how much of an at-risk nuclear power plant's generation should be exempt, Treasury should target a share that is only as large as is absolutely necessary to keep the facility operational, with as much power as possible still serving the grid. However, this share should remain exempt for the 10-year life of the tax credit to provide investment certainty for the electrolyzer. The at-risk facility can return on a periodic basis to request a larger exemption if the market conditions for its remaining power continue to worsen. Treasury could work with the DOE to design the appropriate boundaries for this exemption—such as the frequency of application reviews and any caps on how much of a nuclear facility's output can go toward subsidized hydrogen production—so long as earlier conditions are met that require all other avenues of financial support to be exhausted first.

Formulaic approach

A few qualities subject nuclear power plants to greater risk of retirement relative to other clean energy resources. First, unlike wind, solar, hydropower, or geothermal, nuclear plants must buy fuel in the form of enriched uranium; combined with needing a team of technicians to safely operate the plant, this means that nuclear plants face material marginal operating costs. Second, nuclear facilities are highly inflexible, meaning they have to operate at a loss when power prices fall below their marginal costs or act in advance to ramp down—itself a costly process. Third, nuclear plants are old, at an average age of 42 years, increasing the costs of investing in their continued safe operation.[18]

Most policy support for nuclear power plants in competitive markets has thus been designed to make them "whole" in low power price periods to minimize or eliminate losses in these hours. But it also means an exemption allowing some share of a plant's power to support subsidized electrolysis would only help the plant as a whole to the degree these profits offset losses in the power market on an annual average basis. This can make it difficult to accurately predict how large such an exemption must be to keep the entire plant viable, which might lead Treasury or the DOE to overestimate the share of a facility's power it needs to exempt (as underestimating this share risks causing retirement).

We therefore propose a formulaic approach that would establish an LMP threshold below which a nuclear power plant could exempt its generation from the incrementality requirement. Using a methodology developed by the DOE, this LMP could vary by facility to account for a market's capacity prices (if applicable), state subsidies, and the availability of the 45U tax credit.[xi] The

---

[xi] This approach ought to only kick in after the CNC program's funding has been exhausted, given its intent is to keep at-risk nuclear power plants operational.

threshold LMP would ideally delineate the energy market price above which a nuclear facility is profitable and below which it would need to operate at a loss or ramp down. In these latter hours, the plant would be permitted to redirect as much of its generation as it desires to its own co-located electrolyzers, which would earn the full 45V credit, or sell EACs exempt from incrementality to other hydrogen producers in the same deliverability region.

This action would drive new load without providing new supply, and the result would most often be a fossil fuel power plant ramping up to fill in the gap. However, the relevant counterfactual is one in which the nuclear power plant would otherwise retire, in which case fossil fuel generation would already be higher. Thus, this formulaic approach is beneficial from an emissions perspective because it ensures a nuclear plant's generation fully flows to the grid in higher-value hours (above the threshold LMP) and still partially supports the grid in other hours (depending on how much it redirects to electrolysis, whether directly to a co-located project or via EAC sales). In theory, it also uses market information to ensure facilities only draw on 45V support as much as necessary to remain viable while providing an effective floor for their revenues.

One benefit to this approach is that it encourages investment in flexible electrolyzer technologies and midstream infrastructure like hydrogen storage. This is because, while an application-based approach sets aside a portion of a nuclear plant's generation to feed an electrolyzer around the clock, this formulaic approach results in a varying share of qualifying output on an hour-to-hour basis. Thus, electrolyzers taking advantage of this exemption would need to flexibly ramp up and down to only capture hours in which prices are below the threshold LMP. This arrangement will also be beneficial after 45V expires, as co-located electrolyzers (or at least those located in the same congestion pocket) will continue to act as an effective price floor for nuclear generation, paying up to some price threshold for whatever the plant can offer during lower-LMP hours.

The core challenge to this approach is that finding the threshold LMP that is just right for each nuclear power plant could be a complicated and uncertain process. If the threshold is too low, the at-risk facility won't be able to generate enough revenue from its EACs to stay afloat. If the threshold is too high, then the counterfactual falls apart—some share of the nuclear plant's output would be earning subsidies for electricity it would have sent to the grid regardless, and electrolyzers would be consuming electricity that is actually being generated by fossil fuels. The DOE would need to weigh a range of factors to arrive at an appropriate estimate,[xii] and given these uncertainties, it would be justified in reevaluating these thresholds on a periodic basis.

[xii] For example, the DOE would need to consider all the ways nuclear facilities might use access to exempt EACs, such as (1) putting them all toward building and operating on-site electrolyzers; (2) having them supplement incremental clean energy resources in building and operating on-site electrolyzers; or (3) selling them to other hydrogen producers within the same deliverability region.

Risks and downsides

Despite the options above, Treasury does not need to make an exemption intended to prevent the retirement of nuclear power plants. The risks of inducing GHG emissions beyond 45V's statutory thresholds—or the efforts required to sufficiently mitigate them—are likely not necessary to accept nor worth the limited reward. Four arguments support this conclusion.

First, it's not clear that any U.S. nuclear power plants need access to 45V to stay afloat during the life of the credit. At the time of this paper's publication, there has been no taker for the CNC program's second award cycle.[19] The Palisades nuclear plant in Michigan is even working with a separate program through the DOE's Loans Program Office to secure a loan that would bring it out of retirement.[20] So the combination of today's federal and state policy support appears to be working. The risk of an exemption is that it could worsen emissions, such as if either methodology far overestimates what is actually needed to keep nuclear power plants in the market and causes fossil fuel power plants to ramp up. While the formulaic approach is meant as an elegant, market-based solution, it could still be a challenging, ongoing task for the DOE to administer.

Second, this exemption brings its own market manipulation risks. For example, the owner of a nuclear generation fleet could build new assets in the same local areas, suppressing LMPs and worsening the nuclear plants' economics. This would be an economically inefficient outcome driven by a desire to qualify a greater share of nuclear generation for 45V—a highly attractive subsidy that represents risk-free value for nuclear facilities relative to their reported marginal costs (e.g., $80/MWh revenues from 45V relative to $31/MWh average total generating costs in 2022).[xiii]

Third, while the IRA does not explicitly exclude nuclear power plants from earning both the 45U and 45V tax credits, Treasury does not need to establish an exemption for existing nuclear generation to qualify under 45V for a facility to achieve this outcome. Treasury's draft guidance allows the incremental capacity of uprated facilities to qualify a pro-rated portion of their generation for 45V credits. Thus, an existing nuclear facility that experiences sub-optimal market conditions (thereby qualifying it for 45U) and that undergoes a capacity uprate (thereby qualifying a portion of its output for 45V) would earn tax credits from both provisions.

Fourth, while this may be outside of what Treasury can consider in its rulemaking, an exemption that qualifies existing nuclear generation for 45V is not particularly cost-effective policy. The avoided GHG emissions cost for 45V—assuming nuclear power plants would actually otherwise

---

[xiii] $80/MWh assumes an electrolyzer efficiency of 50 kWh/kg and a hydrogen sale price of $1/kg. The average total generating cost number comes from the Nuclear Energy Institute, though it does not "represent the full costs of operations" such as "market and operational risk management, property taxes, depreciation and interest costs, spent fuel storage costs, or returns on investment." See: https://www.nei.org/CorporateSite/media/filefolder/resources/reports-and-briefs/2023-Costs-in-Context.pdf.

retire—is on the order of $300 per ton of avoided CO_2.[xiv] By contrast, state zero-emission credit programs are highly cost-effective, with The Brattle Group finding a cost of $12 to $20 per ton of avoided CO_2 (and actual costs that may be "considerably lower than this").[21] This makes the 45V credit roughly 15 to 30 times more expensive than programs that top off nuclear power plants' revenues to keep them on the grid. Allowing an exemption for 45V might discourage states or the federal government from providing far more cost-effective nuclear support programs. This prospect is particularly concerning with respect to states, which could shift the cost burden to the federal government by not renewing their support

Takeaways and recommendations

Our assessment of potential exemptions for retirement preventions can be summarized as follows:

- Treasury should not offer an exemption intended to support clean energy resources that might retire if not for support from 45V, as there is little evidence this is needed, and any attempt to capture such circumstances is far more likely to induce substantial GHG emissions (by overestimating facilities' revenue requirements) and market manipulation.
- If Treasury decides to offer an exemption for retirement preventions, two options can help mitigate adverse impacts: an application-based approach and a formulaic approach.
    - An application-based approach would require nuclear facilities to claim CNC program funding, 45U tax credits, and any applicable state program funds to the fullest extent available, then demonstrate via a CNC-style application that additional support is needed to remain financially viable. Even then, the DOE should keep the qualifying generation as low as possible, though this allowance ought to remain for a full 10 years.
    - A formulaic approach would allow nuclear power plants to divert generation to electrolyzers any time their LMP falls below a facility-specific, DOE-determined threshold price that accounts for all other revenue sources (including from capacity markets, state subsidies, 45U, and the CNC program).
- While the IRA does not explicitly bar nuclear facilities from earning both the 45V and 45U credits, withholding a retirement prevention exemption does not close the door to stacking these incentives. Nuclear generators can still undergo a capacity uprate and qualify some share of their output for 45V while earning 45U credits on the rest of their production.

[xiv] This estimate assumes electrolytic hydrogen displaces gray hydrogen downstream for an avoided 10 $kgCO_2/kgH_2$. The cost could rise quickly if an exemption allowed some nuclear generation to qualify for 45V that is not strictly needed to keep the facility operational, or if hydrogen is used in low-value applications like replacing natural gas combustion in furnaces, industrial equipment, or power plants.

BINDING EMISSIONS CAPS

In our original paper on 45V, we do not comment on binding emissions caps outright, except in the context of referencing the EU's rules for the "additionality" principle.[22] Specifically, we note that the EU defers additionality until 2028 and hourly matching (with no grandfathering) until 2030, but its GHG emissions cap makes this phase-in permissible, contrasting this with the U.S. which has no comparable national program.

Below, we discuss what's required to deem an emissions cap as "binding," provide examples from U.S. programs, acknowledge other potential adverse impacts, and note differences from the EU's approach. We conclude that Treasury should only exempt states or regions from the incrementality requirement if it can determine the GHG emissions caps are sufficiently binding.

Principles for judging whether an emissions cap is binding

In theory, a binding emissions cap exemption from incrementality makes sense because it suggests emissions induced by electrolyzers would be offset elsewhere. However, the entire basis of such an exemption rests on how "binding" is defined and determined in the context of Treasury rules. There are at least three sets of questions Treasury ought to ask when making this determination:

- Are the program's **enforcement mechanisms** stronger than the value of the 45V credit? In other words, if an entity can opt out of the need to purchase and retire emissions allowances by paying a fee, is this fee sufficiently high to deter electrolyzer operations whose GHG emissions could not be offset elsewhere? Could electrolyzers freely soak up surplus allowances, thereby raising GHG emissions against 45V's intent (despite being a permissible outcome under a cap-and-trade program's design)?
- Does the program protect against **emissions leakage** from electrolyzer operations within the capped state or region inducing GHG emissions outside of this state or region? Do electricity imports into the covered state or region have their GHG emissions impact properly assessed? Does the program account for electricity losses from transmission and storage (i.e., does the cap affect electricity generation rather than sales)?
- While this may be outside what Treasury can consider, has the program demonstrated **political resilience** to price shocks? Could an influx of heavily-subsidized electrolyzers raise emissions allowance prices and weaken the durability of the entire program, such as by leading policymakers to increase or repeal the cap? On the other hand, might the presence of a binding emissions cap exemption lead other states to tighten or pass their own binding cap-and-trade programs in order to facilitate easier access to the 45V credit?

Ultimately, if electrolyzers can induce GHG emissions despite existing within a carbon cap-and-trade program, Treasury should not exempt them from the incrementality requirement. In this paper, we do not try to make a final judgment on whether states' current programs ought to be considered "binding" in their current form. Instead, we wish to highlight how programs that ostensibly control against rising GHG emissions may not completely do so in reality, and we recommend Treasury consider the above questions when deciding how and whether to design such an exemption.

Examples from existing programs

California

As an example of the “enforcement mechanism” question set, consider California’s cap-and-trade program, which has an Allowance Price Containment Reserve (or Reserve for short). When allowance prices surpass certain thresholds, the California Air Resources Board—which administers the program—will sell limited numbers of additional allowances at $56.20 (the Tier 1 Price in 2024) and $72.21 (the Tier 2 Price in 2024) per metric ton of CO_2. If all allowances are exhausted, any entities that can’t cover their emissions can pay a fee of $88.22 (the Price Ceiling Sale Price in 2024) to comply with the program.[23]

Now consider an electrolyzer in California that is exempt from the incrementality requirement and earns the $3/kg$H_2$ tax credit. The electrolyzer would induce upstream GHG emissions from the marginal natural gas power plant, which we assign as 22.45 kgCO_2/kgH_2 referring to our earlier analysis using the EPA AVERT tool.[24] In this case, the willingness to pay for carbon pollution rights in order to earn the highest-value 45V credit would be up to approximately $134 per metric ton of CO_2. This is 1.5 times as much as California’s Price Ceiling Sale Price and more than twice the Tier 1 Price for unlocking extra emissions allowances.

If the induced emissions from hydrogen production are lower, then the willingness to pay for carbon allowances to capture the 45V credit would rise accordingly. For example, an emissions rate of 10 kgCO_2/kgH_2 would make the electrolyzer effectively willing to pay up to $300 per metric ton CO_2. The opposite would hold true for higher emissions intensities.

In any case, unmitigated incremental demand for hydrogen production (i.e., electrolyzers using exempt EACs from existing clean energy resources and causing natural gas power plants to ramp up) will create new demand for emissions allowances with a willingness to pay well above the cost mitigation thresholds of the existing cap-and-trade program. This means electrolyzers in California could functionally comply with the cap-and-trade program while still contributing to higher levels of GHG emissions—in violation of 45V’s emissions threshold. In other words, the California program is more of a “soft” cap than a binding one—the emissions budget “expands or contracts in response to price bounds set by the legislature and [California Air Resources Board].”[25]

Regional Greenhouse Gas Initiative

As an example of all three question sets, consider the Regional Greenhouse Gas Initiative (RGGI), which is a cap-and-trade program covering the electric sectors of states that opt in (with 11 states in the Northeast and Mid-Atlantic participating as of 2023). RGGI similarly has a Cost Containment Reserve that unlocks additional allowances above a certain price, but its threshold is much lower than California’s at $15.92 per metric ton CO_2 in 2024.[26]

Evidence also suggests RGGI suffers from an emissions leakage problem, with one paper finding 43 to 86 percent of the emissions reduction in RGGI states is a result of leakage, with unregulated areas raising their emissions accordingly.[27] For example, particularly since RGGI only covers a subset of a major U.S. organized power market (PJM), fossil fuel power plants may ramp down in regulated PJM states like Delaware, Maryland, New Jersey, and Virginia while ramping up in unregulated PJM states like Ohio, Pennsylvania, and West Virginia.[xv] RGGI also only covers generators of at least 25 MW in size, so larger generators' reduced output could be made up in part by smaller generators' increased output without completely offsetting GHG emissions.[28] This means even if some of an electrolyzer's GHG emissions are offset as a result of RGGI, a sizable portion may remain, thereby violating 45V's emissions threshold.

Lastly, RGGI has a history of states changing their participation status, meaning there's a risk that rising allowance prices could alter which states remain in the program. For example, New Jersey exited RGGI in 2012 only to rejoin in 2020, and the participation of Virginia (trying to exit) and Pennsylvania (trying to join) is subject to ongoing legal challenges.[29] This dynamic may be beyond the scope of what Treasury can consider, but the reality remains that its decision on how electrolyzers can earn 45V tax credits could have broader implications for the design or longevity of cap-and-trade programs.

Notably, these considerations do not mean cap-and-trade programs are generally ineffective. Instead, they imply such programs may not currently be sufficiently robust to withstand the influence of a highly lucrative federal subsidy that supports an energy-intensive technology like electrolyzers—and thus, Treasury should act with great care in designing any exemption for states or regions with emissions caps, with special attention paid to whether they are truly binding.

Additional considerations

There are a few other ways an incrementality exemption for states or regions with binding emissions caps could have secondary adverse impacts.

First, an exemption could further stress interconnection queues by adding heavily subsidized demands that, by extension, will need new clean supply to make continued progress down an emissions trajectory. For example, CAISO faced an approximate 200 GW backlog in 2022.[30] The incrementality requirement helps ensure supply and demand are brought on together, including via co-located project configurations that can lessen or avoid interconnection challenges (e.g., if built off grid, or if including a grid connection that equals the renewables' capacity minus the electrolyzer's capacity). It also prevents electrolyzers from siphoning power from existing baseload

[xv] Even covering an entire power market may not be enough to prevent leakage given markets can trade amongst each other (e.g., PJM with MISO); however, inter-market transmission connectivity is generally less than intra-market connectivity.

or dispatchable sources of clean energy like nuclear, geothermal, and hydropower, which are very difficult to build and provide much-needed grid reliability services.

Second, in designing the emissions reduction plans that informed an emissions cap policy, states might not have anticipated the extra load they would incur from a 45V exemption. For example, the California Air Resources Board's latest Scoping Plan assumes all hydrogen production comes from either off-grid solar or not from electrolysis.[31] Exempting California projects from the incrementality requirement would almost certainly put these loads on the grid, making California's targets tougher to achieve.

Third, an exemption could end up building the type of clean hydrogen industry that is poorly positioned to survive after 45V expires. Electrolyzers paired with existing nuclear, geothermal, or hydropower could run around-the-clock—or at least have ramping periods that are gradual and predictable far in advance—even under an hourly matching requirement. As a result, developers would most likely build the cheapest, least-flexible electrolyzers to maximize their profit under 45V. An exemption from hourly matching would make matters worse by extending the issue to wind and solar as well. These projects would be ill-equipped to adjust to a post-45V world, nor would they have helped make flexible technologies less expensive.

Separately, we note that the existence of a 100 percent clean electricity standard (CES) should generally not allow for an exemption from incrementality. Without incrementality, a state would quickly add new loads that would only be matched with clean energy up to the current interim CES requirement. That is, at a 50 percent target, developers would only need half of an electrolyzer's load to be matched with clean energy, meaning the rest would be dirty. Such a 50-50 split of renewables and natural gas would result in an average hydrogen production emissions rate of roughly 10 $kgCO_2/kgH_2$—more than 20 times the IRA's requirement for earning the $3/kg credit and equivalent with steam methane reformation. Even when a state has realized its 100 percent CES target, incrementality would be a necessary (if not sufficient) feature of maintaining this achievement, meaning an exemption has no value.

Differences between the U.S. and EU

The EU has an emissions cap, and we note in our prior 45V paper that this contributes to it being reasonable for the EU to have a more lenient phase-in period for its incrementality and hourly time-matching requirements.[32] However, it's worth further exploring the differences between U.S. and EU emissions caps and approaches to qualifying hydrogen as "clean."

First, the EU cap is geographically much broader than U.S. state and regional programs, meaning it suffers less from emissions leakage across borders. If the U.S. adopted a nation-wide emissions cap, an exemption would hold more validity.

Second, the EU's incrementality requirement is more stringent than that of Treasury's proposal. Both require electrolyzers to draw power from clean energy resources that came online no later than 36 months before the electrolyzers' commercial operations. However, only the EU rules require long-term power purchase agreements and prevent the contracted clean energy resources from earning other subsidies. This means a temporary incrementality exemption (offered until 2028) would do less harm if the EU's emissions cap isn't fully binding or air-tight from leakage.

Third, the EU does eventually require compliance with all three pillars regardless of the existence of its emissions cap. It thus recognizes that its cap may not be truly binding or that there are other reasons why such requirements may be beneficial for the clean hydrogen industry's growth and other outcomes like electric system reliability.

Thus, in considering whether and how to design an exemption for binding emissions caps, it's worth Treasury exploring how the EU program differs from U.S. state or regional programs as well as why the EU requires eventual compliance with all three pillars despite its cap.

Takeaways and recommendations

Our assessment of potential exemptions for states and regions with binding emissions caps can be summarized as follows:

- If Treasury decides to offer an exemption for states or regions with binding emissions caps, the exemption should be limited to the incrementality requirement, and Treasury should set strong guardrails around what counts as "binding." These guardrails should cover:
    - Whether enforcement mechanisms are sufficiently strong to withstand 45V's high incentive value without opening the door to additional GHG emissions allowances, as well as whether the program has a glut of allowances that electrolyzers could consume without GHG emissions needing to be offset elsewhere;
    - Whether the program sufficiently controls for emissions leakage into unregulated regions or sectors, such that GHG emissions induced by an electrolyzer's operations are offset in full rather than merely in part, (with such a determination likely requiring tighter deliverability regions and independent modeling); and
    - Whether the program is resilient to shocks in the price of emissions allowances such that an exemption would not contribute to policymakers weakening or repealing the emissions cap.
- While it's possible that no states or regions would currently qualify for such an exemption, these jurisdictions would understand what steps they need to take to strengthen their emissions cap programs to meet an appropriate definition of "binding," giving them agency to make the decision that is in their best interest.
- Regardless of Treasury's approach, risks remain from offering such an exemption, including exacerbating interconnection queue backlogs, upending state emissions scoping plans, and watering down 45V's incentive to invest in flexible electrolyzer technologies.
- The EU has a geographically broader emissions cap than any U.S. program, its rules for clean electrolytic hydrogen include tougher requirements on incrementality, and it still

ultimately requires all three pillars despite having a cap; thus, it's worth Treasury considering what it can learn from the EU's program design and its decision to continue requiring compliance with a full three-pillars framework.

CONCLUSION

With its core requirements of incrementality, deliverability, and time-matching, Treasury struck the right balance between the accuracy with which it can verify electrolyzers' GHG emissions and the rules' administrability. However, in an understandable attempt to qualify a few edge cases that otherwise are excluded in a three-pillars framework, Treasury risks poisoning the well. Namely, a general carve-out for existing clean energy—as well as targeted approaches that don't include sufficiently careful guardrails—could open the door to a substantial amount of dirty hydrogen qualifying for the highest-value tax credit by way of loopholes. Such exemptions could also dilute the incentive to innovate on flexibility, harming the clean hydrogen industry's long-term prospects. These allowances may seem small (e.g., a mere 5 percent of existing clean energy), but their adverse impacts could punch well above their weight.

We recommend Treasury preserve its three-pillars system and exclude any general carve-out. To the extent Treasury pursues targeted exemptions, it should design them in a manner that does not toss accuracy to the wind in favor of ease of administrability—a goal we think is achievable (if not necessary) for curtailment and binding emissions caps but likely out of reach for retirement preventions. By maintaining this important balance between accuracy and administrability, Treasury can ensure 45V builds a truly clean hydrogen industry capable of thriving long after the credit expires.

REFERENCES

[1] "U.S. Department of the Treasury, IRS Release Guidance on Hydrogen Production Credit to Drive American Innovation and Strengthen Energy Security," U.S. Department of the Treasury, January 25, 2024, https://home.treasury.gov/news/press-releases/jy2010.

[2] "Smart Design Of 45V Hydrogen Production Tax Credit Will Reduce Emissions And Grow the Industry," Energy Innovation: Policy and Technology LLC, April 2023, https://energyinnovation.org/publication/smart-design-of-45v-hydrogen-production-tax-credit-will-reduce-emissions-and-grow-the-industry/.

[3] Energy Innovation, "Smart Design of 45V Hydrogen Production Tax Credit."

[4] U.S. Department of Energy, "Assessing Lifecycle Greenhouse Gas Emissions Associated with Electricity Use for the Section 45V Clean Hydrogen Production Tax Credit," December 2023, https://www.energy.gov/sites/default/files/2023-12/Assessing_Lifecycle_Greenhouse_Gas_Emissions_Associated_with_Electricity_Use_for_the_Section_45V_Clean_Hydrogen_Production_Tax_Credit.pdf.

[5] U.S. Environmental Protection Agency, "45V NPRM EPA Letter," December 20, 2023, https://home.treasury.gov/system/files/136/45V-NPRM-EPA-letter.pdf.

[6] U.S. Department of the Treasury, "Section 45V Credit for Production of Clean Hydrogen; Section 48(a)(15) Election to Treat Clean Hydrogen Production Facilities as Energy Property" (U.S. Department of the Treasury, December 22, 2023), 37, https://public-inspection.federalregister.gov/2023-28359.pdf.

[7] "Section 45V Credit for Production of Clean Hydrogen; Section 48(a)(15) Election to Treat Clean Hydrogen Production Facilities as Energy Property," 42.

[8] "Section 45V Credit for Production of Clean Hydrogen; Section 48(a)(15) Election to Treat Clean Hydrogen Production Facilities as Energy Property," 45.

[9] "Curtailment: When We Throw Away Clean Energy," GridStatus.io, November 7, 2023, https://blog.gridstatus.io/curtailment/.

[10] Phillippe K. Phanivong, "California's Growing Solar and Wind Problem," EcoBlock, December 5, 2023, https://ecoblock.berkeley.edu/blog/californias-growing-solar-and-wind-problem/.

[11] "ICAP ETS Map," International Carbon Action Partnership, accessed February 19, 2024, https://icapcarbonaction.com/en/ets.

[12] Energy Innovation, "Smart Design of 45V Hydrogen Production Tax Credit."

[13] Ben King et al., "How Clean Will US Hydrogen Get? Unpacking Treasury's Proposed 45V Tax Credit Guidance," *Rhodium Group* (blog), January 4, 2024, https://rhg.com/research/clean-hydrogen-45v-tax-guidance/.

[14] Energy Innovation, "Smart Design of 45V Hydrogen Production Tax Credit."

[15] Energy Innovation, "Smart Design of 45V Hydrogen Production Tax Credit."

[16] GridStatus.io, "Curtailment."

[17] Energy Innovation, "Smart Design of 45V Hydrogen Production Tax Credit."

[18] U.S. Energy Information Administration, "U.S. Nuclear Industry," August 24, 2023, https://www.eia.gov/energyexplained/nuclear/us-nuclear-industry.php.

[19] Brian Dabbs, "No Takers for Biden's Nuclear Bailout," E&E News, January 8, 2024, https://subscriber.politicopro.com/article/eenews/2024/01/08/no-takers-for-bidens-nuclear-bailout-00134067.

[20] Timothy Gardner, "Holtec to Get $1.5 Bln Loan to Re-Open Michigan Nuclear Power Plant -Source," *Reuters*, January 31, 2024, https://www.reuters.com/business/energy/holtec-get-15-bln-loan-re-open-michigan-nuclear-power-plant-source-2024-01-31/.

[21] Metin Celebi et al., "Nuclear Retirement Effects on CO_2 Emissions: Preserving a Critical Clean Resource" (The Brattle Group, December 2016), https://www.brattle.com/wp-content/uploads/2017/10/5727_brattle_nuclear-carbon_whitepaper_-_dec2016.pdf.

[22] "Smart Design Of 45V Hydrogen Production Tax Credit Will Reduce Emissions And Grow the Industry."

[23] "Cost Containment Information," California Air Resources Board, accessed February 19, 2024, https://ww2.arb.ca.gov/our-work/programs/cap-and-trade-program/cost-containment-information.
[24] "Smart Design Of 45V Hydrogen Production Tax Credit Will Reduce Emissions And Grow the Industry," app. A.
[25] James Bushnell, "California's Cap-and-Trade Market Enters Its Teen-Age Years," *Energy Institute Blog, UC Berkeley* (blog), November 27, 2023, https://energyathaas.wordpress.com/2023/11/27/californias-cap-and-trade-market-enters-its-teen-age-years/.
[26] "Elements of RGGI," RGGI, accessed February 19, 2024, https://www.rggi.org/program-overview-and-design/elements#.
[27] Yishu Zhou and Ling Huang, "How Regional Policies Reduce Carbon Emissions in Electricity Markets: Fuel Switching or Emission Leakage," *Energy Economics* 97 (May 1, 2021): 105209, https://doi.org/10.1016/j.eneco.2021.105209.
[28] "Elements of RGGI."
[29] "Regional Greenhouse Gas Initiative (RGGI)," *Center for Climate and Energy Solutions* (blog), accessed February 19, 2024, https://www.c2es.org/content/regional-greenhouse-gas-initiative-rggi/.
[30] Joseph Rand et al., "Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2020" (Lawrence Berkeley National Laboratory, April 2023), https://doi.org/10.2172/1784303.
[31] "2022 Scoping Plan for Achieving Carbon Neutrality" (California Air Resources Board, December 2022), 202, https://ww2.arb.ca.gov/sites/default/files/2023-04/2022-sp.pdf.
[32] "Smart Design Of 45V Hydrogen Production Tax Credit Will Reduce Emissions And Grow the Industry."