

TECHNOLOGY-NEUTRAL VERSUS TECHNOLOGY-SPECIFIC PROCUREMENT*

Natalia Fabra and Juan-Pablo Montero

An imperfectly informed regulator needs to procure multiple units of some good (e.g., green energy, market liquidity, pollution reduction, land conservation) that can be produced with heterogeneous technologies at various costs. How should she optimally procure these units? Should she run technology-specific or technology-neutral auctions? Should she allow for partial separation across technologies? Should she instead post separate prices for each technology? What are the trade-offs involved? We find that one size does not fit all: the preferred instrument depends on the costs of the available technologies, their degree of substitutability, the extent of information asymmetry and the costs of public funds. We illustrate the use of our theory for policy analysis with an ex ante evaluation of Spain's recent renewable auction.

Spain has introduced a novel auction design to procure renewable energy: a joint auction for solar and wind, but with minimum quotas reserved for each technology.¹ Despite the novelty of this design, Spain is just one of many countries resorting to renewable energy auctions to reduce carbon emissions at the lowest possible fiscal cost. According to IRENA (2019), by the end of 2018 more than 100 countries had used auctions to procure renewable energy, i.e., a ten-fold increase in just one decade.² As Fowlie (2017) puts it, a worldwide ‘renewable-energy-auction revolution’ is underway. Remarkable in this revolution is the fact that no two auction designs look alike. They often differ in several dimensions, ranging from the pricing format to the contract duration, to name just two.

One key dimension, which is the focus of this paper, is whether auction designs are technology neutral, or whether they discriminate across technologies, either by type, location and/or scale.³

* Corresponding author: Juan-Pablo Montero, Department of Economics, Pontificia Universidad Católica de Chile, Vicuña Mackenna 4860, Santiago, Chile. Email: jmontero@uc.cl

This paper was received on 14 July 2021 and accepted on 29 September 2022. The Editor was Heski Bar-Isaac.

The data and codes for this paper are available on the Journal repository. They were checked for their ability to reproduce the results presented in the paper. The replication package for this paper is available at the following address: <https://doi.org/10.5281/zenodo.7087620>.

This project has received funding from the European Research Council (ERC) under the European Union Horizon 2020 Research and Innovation Program (Grant Agreement No. 772331). Montero also acknowledges support from Fondecyt (Grant No. 1210252), and the Chair in Energy and Prosperity at Ecole Polytechnique. We have benefited from comments by the Editor, Heski Bar-Isaac, and three anonymous referees, as well as by David Andrés-Cerezo, Peter Cramton, Peter Eccles, Vitali Gretschko, Tibor Heumann, Dave Kelly, Paul Klemperer, Leslie Marx, Guy Meunier, Mateus Souza and seminar participants at the Energy Workshop (Universidad Carlos III-TSE), SURED Conference (ETH Ancona), Université de Montpellier, University of Miami, TILEC (Tilburg), RIDGE Environmental Economics Workshop, PUC-Chile, INRA-Paris, the Virtual Market Design seminar, WIPE (Reus) and the University of Zurich. Imelda, Vicente Jiménez and Mateus Souza provided excellent research assistance.

¹ This design was first used in January 2021. The resulting prices have been highly competitive according to all international standards (IRENA, 2020). Two-thirds of the total auctioned volume have been allocated to solar projects, just before triggering the minimum quota reserved for wind. It has also been used in October 2021, also resulting in a competitive outcome.

² Furthermore, many large corporations are also resorting to auctions to procure renewable power. For instance, from 2017–9 Google purchased a number of renewables equivalent to 100% of the company's total electricity use (Google, 2020).

³ The European Commission's (2014) Guidelines on State aid for environmental protection and energy (currently under revision) require that auction schemes treat all technologies on a non-discriminatory basis (technology neutral),

Yet other auctions rely on hybrid designs that allow for some degree of competition across technologies while favouring some over others, e.g., by giving a handicap to some technologies, or by guaranteeing them a minimum quantity allocation, as recently done in Spain. Why is there such a large variation in auction designs regarding the treatment of the various technologies? What are the trade-offs involved? Is it possible to identify a technology approach that performs better than the formats currently in use? The objective of this paper is to provide a sufficiently general framework to understand, from a purely economic regulatory perspective, how to optimally procure the various technologies, and when and why a particular procurement approach should be preferred over another.

Beyond green energy, the question of how to procure goods or services in the presence of multiple technologies is relevant in a wide variety of public procurement settings. Another notable example arises in the context of the liquidity auctions ran by central banks, in which borrowers (i.e., commercial banks) offer either strong or weak collateral in exchange for liquidity (Klemperer, 2010; Frost *et al.*, 2015). In the past, central banks have considered different options, from posting prices (i.e., interest rates), to running separate auctions for each type of collateral, to running a joint auction for both types of collateral.⁴ The choice between these approaches is also relevant in settings such as procurement of pollution reductions (Laffont and Tirole, 1996; van Benthem and Kerr, 2013) and land conservation (Mason and Plantinga, 2013), among others.

In this paper, we develop a simple model to identify and properly weigh the key factors involved in technology procurement design in practice. We consider two types of technologies, say, solar and wind,⁵ and a continuum of suppliers of each technology.⁶ We capture the regulator's incomplete information by assuming that supply curves are subject to positively or negatively correlated shocks across technologies. The regulator's objective is to maximise (expected) social benefits minus total costs, subject to a budget constraint that gives rise to costly public funds. In solving the regulator's problem, we restrict attention to procurement formats that rely on uniform pricing.⁷

We start our analysis by showing that the optimal mechanism is a product-mix auction *à la* Klemperer (2010), i.e., a single auction where the regulator commits to a demand schedule that is contingent on the bids submitted for the two technologies. This allows the total quantity procured, as well as the quantity allocated to each technology, to adjust to the cost shocks. Furthermore,

with only a few exceptions allowed. This has prompted a shift for which the number of technology-neutral auctions in Europe increased from one in 2015 to 18 in 2019 (Jones and Pakalkaite, 2019). Still, there exist many technology- or location-specific mechanisms in place. For instance, the 2009 European Union's Renewable Energy Directive determines renewable targets at the national level, with no trading across countries.

⁴ Some joint auctions have followed the product-mix design proposed by Klemperer (2010) to the Central Bank in the UK, where the auctioneer announces demand schedules for the different types of collateral, and banks have the opportunity to parsimoniously express substitutable preferences over them. We come back to this design shortly.

⁵ Manzano and Vives (2021) also considered a divisible good uniform-price auction with two groups of identical bidders. In their model, bidders compete in demand schedules and do not know their costs. The attempt to learn costs from the market price shapes their bidding behaviour, leading them to submit flatter or steeper demand functions. Another key difference with our model is that their welfare analysis does not incorporate the social cost of public funds.

⁶ Price-taking behaviour not only facilitates the analysis, but also captures, to a large extent, what we have seen in recent renewable auctions (the January 2021 Spanish auction, for example, had 84 different bidders, offering more than three times the auctioned amount, with a final number of 32 winners). Similarly, Lamp *et al.* (2021) showed that German renewable auctions have been highly competitive. In any event, in Appendix C we discuss to what extent market power may change some of our results.

⁷ In practice, regulators have used both uniform-price and pay-as-bid formats. For instance, the German renewable auctions had a uniform-price format until 2015, and a pay-as-bid format thereafter (Lamp *et al.*, 2021). Previous papers have shown that, in the context of our analysis (i.e., perfectly competitive auctions), the two formats give rise to the same outcome (Fabra *et al.*, 2006). Under strategic behaviour, equilibrium outcomes might differ across formats, but only in the presence of pivotal bidders (i.e., the capacity of all bidders is needed to cover total demand).

whenever the regulator cares about payments (i.e., public funds are costly), the quantities allocated across technologies depart (ex post, but also ex ante) from the cost-minimising solution even when the two technologies are perfect substitutes. This gives rise to different prices for the two technologies, despite their benefits being the same—a result that is reminiscent of third-degree price discrimination (Bulow and Roberts, 1989). In sum, the optimal mechanism strikes a balance between efficiency and rent extraction: it gives up full efficiency in order to reduce procurement costs. Although this rent-efficiency trade-off has been widely recognised in the literature of regulation and public procurement (Laffont and Tirole, 1993; Laffont and Martimort, 2002), its impact on the preferred regulatory instrument to procure multiple technologies has not been systematically analysed before.

While the optimal mechanism allows the regulator to fully overcome her information asymmetry, it has never been used in practice (at least not in the realm of resource and renewable-energy auctions, as our discussion above attests). Instead, regulators often rely on simpler policy designs that adjust only partially to actual cost realisations. Under these simpler mechanisms, regulators cannot escape the rent-efficiency trade-off described above, which is a centrepiece in the rest of our analysis.

Motivated by the renewable auction revolution, we first consider the case of quantity regulation, i.e., procurement auctions. We start with two of the simplest designs found in practice: the regulator has to commit ex ante to procure a given number of units in either a single technology-neutral uniform-price auction or in two technology-specific uniform-price auctions. This lack of flexibility in the total quantity procured implies that these two formats depart from the optimal mechanism. But there are additional reasons that make these two formats sub-optimal.

First, the technology-neutral approach is similar to the optimal mechanism in that the two technologies compete within the same auction, allowing the quantity allocation to adjust to the cost shocks. However, while technology neutrality is effective in minimising costs, it may result in over-compensation. Indeed, because the regulator lacks the ability to discriminate among heterogeneous sources, the regulator may leave too many rents with the more efficient suppliers, unnecessarily increasing procurement costs. Furthermore, the technology-neutral approach may give rise to distorted technology choices, given that it does not internalise the degree of substitutability across technologies.

Second, the technology-specific approach is similar to the optimal mechanism in that both technologies receive different prices, with the allocated quantities departing from a pure-cost ranking. The objective is two-fold: to reduce rents and to capture the benefits from technology substitutability. However, since the regulator has to choose the technology allocation ex ante, without knowledge of the various costs, the technology-specific approach might not only result in inefficient, but also more costly allocations.

It follows that the choice between a technology-neutral and a technology-specific approach again faces the regulator with a trade-off between minimising costs and minimising rents left to firms. Whenever cost minimisation is more important, technology neutrality should be favoured; whenever rent minimisation is more important, the technology-specific approach should be favoured instead. The choice between the two approaches also depends on the degree of substitutability across technologies, as it might distort the technology allocation under technology neutrality, but not under the technology-specific approach. The choice thus has to be assessed on a case-by-case basis, depending on the values of relevant parameters.

For instance, a well-informed regulator should always run separate auctions, with the technology-specific targets chosen to balance cost minimisation, rent extraction and technol-

ogy substitution (this replicates the outcome of the optimal mechanism).⁸ A similar prescription should be followed if the two technologies are subject to perfectly correlated shocks: in this case, cost minimisation is not in danger either, but technology separation allows one to reduce rents and capture the benefits from technology substitutability.

As incomplete information mounts, however, minimising costs through technology separation becomes increasingly challenging as quantity targets do not adjust to the cost shocks. Eventually, technology neutrality may dominate technology separation unless the costs for the regulator of not discriminating technologies are too large. This ultimately depends on the amount of over-compensation to the more efficient suppliers—as captured by the expected cost difference across technologies—and the unit price of this over-compensation—which depends on the shadow cost of public funds as well as on the degree of substitutability across technologies.⁹ The higher the degree of substitutability across the two technologies, the more likely it is that technology neutrality may dominate, all else equal.

Since neither technology neutrality nor technology separation succeeds in containing both costs and payments, one may argue in favour of hybrid approaches that allow for some partial separation between technologies. Indeed, a handful of countries currently rely on a partial separation approach referred to as ‘technology banding’ for setting renewable support. The idea is to run a single uniform-price auction with suppliers of the *ex ante* inefficient technology (or less resourceful location) receiving a handicap to compete more effectively with suppliers of the *ex ante* more efficient technology or location (Myerson, 1981).¹⁰

Whereas one may speculate that the banding approach is superior relative to the two extremes of full neutrality or full separation, this is not always the case. Trivially, banding dominates technology neutrality as one can always set a neutral handicap. However, through banding one cannot replicate the same outcome as under technology separation. Indeed, we find that banding does not always dominate the technology-specific approach. Not only is the latter better equipped at containing total payments, but more surprisingly, it might also lead to lower costs as compared to banding. The problem with banding is that the handicap that is designed to contain payments also distorts technology substitution away from the efficient allocation. Cost shock volatility, coupled with convex costs, implies (due to Jensen’s inequality) that expected costs under banding might be higher than under the technology-specific approach. This is particularly the case when the correlation of cost shocks is sufficiently high.¹¹

Another hybrid approach, used in Spain’s recent renewable auctions, is the establishment of minimum technology quotas (MTQs) in otherwise technology-neutral auctions.¹² Unlike

⁸ In the context of carbon trading across countries, Martimort and Sand-Zantman (2016) also found that preventing trade across countries is part of the optimal mechanism, insofar as it allows to control rents going to the different countries.

⁹ Adding market power to the model brings new insights. Under technology-specific auctions, market power makes it optimal to further distort the quantity targets, giving rise to more productive inefficiency as compared to the technology-neutral approach. While such quantity distortions also allow one to reduce rents, the regulator’s ability to do so through technology separation is diminished the more market power there is. Hence, market power tends to favour the technology-neutral approach. See Appendix C.

¹⁰ Very often, banding is also used to penalise technologies that are considered less valuable, or to incentivise the more valuable ones. For instance, in the renewable auctions in Mexico, plants that have a generation profile that matches the system’s needs receive an additional remuneration, while plants with less valuable production profiles are penalised (IRENA, 2019). Yet, we show that banding can be useful as a payment containment device even in settings in which all technologies are equally valuable.

¹¹ Using a similar framework, but in the context of integrating pollution permit markets, Montero (2001) also found that in some cases a corner solution (alike technology separation in our set-up) may be optimal.

¹² Yet another hybrid option is to introduce technology-specific reserve prices instead of minimum quotas (reserve prices have been used in the auctioning of pollution permits for instance; see Borenstein *et al.*, 2019). As this hybrid

banding, MTQs can be designed to replicate the two extremes of technology neutrality and technology separation. By separating technologies when cost realisations make technologies diverge, MTQs are effective in containing payments when this is most needed. Likewise, by preserving neutrality when cost shocks make technologies more symmetric, MTQs are effective in avoiding cost inefficiencies. However, this does not mean that MTQs are always superior to banding. Indeed, we show that banding may dominate MTQs when one technology is clearly more efficient than the other and their costs are not too positively correlated.

So far, we have considered a regulator who procures a given number of units under different auction formats. Those scenarios can arise when the total quantity to be procured is not under the regulator's control, but rather exogenously given (as in Spain's recent auctions); for instance, in response to a higher-level country commitment to reduce carbon emissions. The case of an endogenous total quantity opens a new set of questions. In particular, it may no longer be preferable to rely on quantity-based instruments (e.g., auctions), but rather on price-based instruments (e.g., feed-in tariffs)—indeed, as we discussed above, the optimal mechanism involves *ex post* quantity adjustments. However, prices also depart from full optimality since they have to be chosen *ex ante*, without any adjustment to cost shocks. To study this additional instrument choice problem, we extend Weitzman (1974) by considering multiple technologies and costly public funds. New insights emerge.

If, on the one hand, technology-specific auctions happen to dominate a technology-neutral auction, the comparison of 'prices versus quantities' gives rise to a modified version of Weitzman's (1974) seminal expression.¹³ In this case, the presence of multiple technologies enhances the superiority of prices over quantities, since the former allows the quantities of the various technologies (and not just the total quantity) to better adjust to cost shocks. If, on the other hand, a technology-neutral auction happens to dominate technology-specific auctions, the comparison of prices versus quantities includes an additional term: a rent-extraction term. When public funds are not too costly, such that the rent-extraction term is not too large, a single quantity target may still dominate two prices, as it allows for more quantity adjustment across technologies.

Motivated by Spain's recent renewable auctions, we close the paper with an application that serves to illustrate the use of our theory for policy analysis. We ask ourselves, from an *ex ante* perspective (i.e., as of 2020), what would be the pros and cons of auctioning 3,000 MW of renewable energy under an MTQ approach, as Spain did in January 2021, relative to other mechanisms, including the optimal one? Since much of the auction data (e.g., actual bids and projects' locations) have not been made available yet, we rely on detailed cost information of the renewable projects that applied for permission during 2019. If the (relative) state of technologies in 2019 is a good proxy for the state of technologies expected for 2021, then our simulations suggest that Spain's novel MTQ design might have been a good (*ex ante*) choice over alternative formats.¹⁴ In fact, a well-designed MTQ format is not far away from implementing the optimal outcome, giving rise to social costs that are only 0.1% to 0.8% higher than under the optimal mechanism. This is in contrast to the other formats, which give rise to a much higher increase

option might result in the total quantity not being fully allocated, we do not cover this case in our analysis. It is tangentially covered when we discuss 'prices versus quantities' in Section 5.

¹³ Our expression coincides with Weitzman's (1974) only when the cost shocks are perfectly correlated and the two technologies are perfect substitutes, as in this case, the two technologies behave just like one.

¹⁴ For completeness, the Spanish design differs in other dimensions not considered in our simulations, so actual and simulated results are not readily comparable. An empirical analysis of the actual auction's outcome is left for future work once bid data become available.

in social costs (up to 12% under technology neutrality, up to 3% under the technology-specific approach or up to 8% under banding).¹⁵

The rest of the paper is organised as follows. Section 1 describes the model. Section 2 characterises the optimal mechanism. Section 3 characterises the solutions under technology-neutral and technology-specific auctions in their simplest formats and compares them to the optimal mechanism both for the cases of perfect and imperfect substitutabilities across technologies. Section 4 analyses and compares two hybrid schemes: technology banding and MTQs. Section 5 analyses price regulation. Section 6 contains the application to Spain's renewables auction. Section 7 concludes. Lengthy proofs, as well as the analysis of market power, are relegated to the Appendices.

1. The Model

There are two types of technologies, say, solar and wind, denoted by 1 and 2. Each technology $t = 1, 2$ can be supplied by a continuum of (risk-neutral) price-taking firms with unit capacity,¹⁶ whose mass is normalised to one.¹⁷ Their (long-run) unit costs are uniformly distributed over the interval $[\underline{c}_t, \bar{c}_t]$, where $\underline{c}_t = c_t + \theta_t$ and $\bar{c}_t = c_t + \theta_t + \gamma$.¹⁸ Therefore, the aggregate cost of supplying $q_t \in [0, 1]$ units of technology t is given by the quadratic function

$$C_t(q_t) = (c_t + \theta_t)q_t + \frac{1}{2}\gamma q_t^2,$$

where $\gamma > 0$ is common to both technologies and $\theta_t \in [\underline{\theta}_t, \bar{\theta}_t]$ is a 'cost shock' that captures the regulator's incomplete information about the costs of supplying technology t (both c_t and γ are public information). We allow c_t and θ_t to differ across technologies. In particular, we assume cost shocks to be jointly distributed according to the PDF $g(\theta_1, \theta_2)$, with $E[\theta_t] = 0$, $E[\theta_t^2] = \sigma_t^2 > 0$ and $E[\theta_1\theta_2] = \rho\sigma_1\sigma_2$, where $\rho \in [-1, 1]$. Thus, we allow cost shocks to be either positively or negatively correlated across technologies.¹⁹

Without loss of generality, we index technologies such that $c_1 \leq c_2$, implying that technology 1 is ex ante more efficient than technology 2. We use $\Delta c \equiv c_2 - c_1 \geq 0$ and $\Delta \theta \equiv \theta_2 - \theta_1$. We further assume that the cost shock intervals $[\underline{\theta}_t, \bar{\theta}_t]$ for $t = 1, 2$ are such that, under all the formats we compare, in equilibrium both technologies get deployed with probability one.²⁰

The deployment of these technologies creates social benefits, which we also capture with a quadratic function of the form

$$B(q_1, q_2) = b(q_1 + q_2) + \frac{\beta}{2}(2 - \eta)(q_1^2 + q_2^2) + \beta\eta q_1 q_2$$

¹⁵ Actual numbers would vary, depending on the regulator's degree of cost uncertainty, but the comparison across formats would remain unchanged.

¹⁶ In Appendix C, we add market power to the analysis.

¹⁷ The same firm could be supplying both technologies. However, since firms are price takers and there are no scope economies across the two technologies, it does not change the analysis whether they offer a single technology or both.

¹⁸ Unit costs are increasing and uncertain, partly because sites vary in quality, as captured in our simulation exercise in Section 6 and also emphasised in Schmalensee (2012).

¹⁹ In some passages we adopt the assumption that $\sigma_1 = \sigma_2 = \sigma$, but only to simplify the exposition.

²⁰ Essentially, this implies that the costs of the two technologies cannot differ too much, either ex post or ex ante. Otherwise, only the low-cost technology would get deployed and there would be no meaningful multi-technology competition. Allowing for this possibility would make the model less tractable without adding new insights.

with $b > 0$ and $\beta < 0$, and where $\eta \in [0, 1]$ captures the degree of substitutability across technologies, from being independent ($\eta = 0$) to perfect substitutes ($\eta = 1$).²¹ Note that, even though both technologies appear symmetric from a benefits standpoint, they are still differentiated from a cost standpoint. The symmetry on the benefit side helps not to bias the analysis in favour of technology-specific approaches. Indeed, if one technology is more valuable than the other, the range of parameter values for which the technology-specific approach would dominate technology neutrality would be obviously enlarged. Furthermore, we assume that b is large enough so that it is always optimal to procure some units.

The risk-neutral regulator's objective is to maximise (expected) social welfare subject to a budget constraint,

$$W(q_1, q_2) = E[B(q_1, q_2) - C(q_1, q_2) - \lambda T(q_1, q_2)], \quad (1)$$

where $C(q_1, q_2)$ denotes the cost of supplying q_1 and q_2 units, $T(q_1, q_2)$ denotes the regulator's total payment and $\lambda \geq 0$ is the shadow cost of public funds (Laffont and Tirole, 1993). We refer to $C(q_1, q_2) + \lambda T(q_1, q_2)$ as the *social cost*, which takes into account both the actual production costs as well as the costs of the fiscal distortions. This formulation is general enough to accommodate different procurement instruments. The functions $C(q_1, q_2)$ and $T(q_1, q_2)$ will take different forms under the various instruments.

The timing of the procurement game is as follows. On date 1, the regulator announces the procurement format and its clearing rules. We restrict attention to formats that rely on uniform pricing, regardless of whether the regulator is using quantity or price schemes.²² On date 2, firms observe the cost shocks θ_1 and θ_2 , and submit their bids. Since truthful bidding is a weakly dominant strategy for a price-taking firm, we adopt cost bidding as the equilibrium strategy (unless explicitly mentioned otherwise). On date 3, prices and quantities are chosen and payments are made.

2. The Optimal Mechanism

We start by characterising the optimal mechanism, which we use as a benchmark to assess the mechanisms that are used in practice. Within the class of mechanisms that rely on uniform pricing, we show that the optimal mechanism is a single auction in which the two technologies are simultaneously sold, possibly at different prices (as in the product-mix design proposed by Klemperer (2010) for Central Banks' liquidity auctions).

LEMMA 1. *The optimal mechanism is a product-mix auction. It is characterised by the regulator's announcement of demand schedules*

$$P_t^d(q_1, q_2) = \frac{\partial B(q_t, q_{-t}) / \partial q_t - \lambda C_t''(q_t) q_t}{1 + \lambda}$$

with firms bidding according to technology-specific supply schedules $P_t^s(q_t)$ for $t = 1, 2$.

²¹ This quadratic formulation is widely used to model multi-product demand systems; see, e.g., Shubik and Levitan (1980). Note also that, when $\eta = 1$, $B(q_1, q_2)$ only depends on the aggregate quantity of the two technologies, $q_1 + q_2$.

²² Note that, since firms face the same cost shock, a pay-as-bid (technology-specific) auction would be equivalent to a uniform-price (technology-specific) auction. Extending this equivalence to a technology-neutral auction would require firms to observe both cost shocks.

PROOF. Given that there is a continuum of price-taking firms, they bid truthfully. Truthful bidding leads to supply schedules given by the marginal cost of each technology, i.e.,

$$P_t^s(q_t; \theta_t) = C_t'(q_t)$$

for $t = 1, 2$. It is then straightforward to show that the resulting prices and quantities, which are obtained from the system $P_t^d(q_1, q_2) = P_t^s(q_t; \theta_t)$ for $t = 1, 2$, solve the same problem of a regulator who observes θ_1 and θ_2 . $\square$

To facilitate the exposition, we first consider the case of perfect substitutes, $\eta = 1$,²³ and leave for later the case of imperfect substitutes, $\eta \in [0, 1)$. Using our quadratic functional forms, the first-best total quantity actually procured can be conveniently decomposed into a deterministic and a stochastic component as (see Appendix A for details)

$$Q^{FB}(\theta_1, \theta_2) = \bar{Q} - \hat{Q}(\theta_1, \theta_2). \quad (2)$$

The stochastic component,

$$\hat{Q}(\theta_1, \theta_2) = \frac{1 + \lambda}{(1 + 2\lambda)\gamma - 2\beta}(\theta_1 + \theta_2),$$

is a function of the cost shocks, allowing the total quantity to adjust ex post. The regulator can perfectly anticipate the optimal quantity only if cost shocks are perfectly and negatively correlated, thus providing a perfect hedge, i.e., if $\theta_1 = -\theta_2$, or when the benefit function is perfectly inelastic, i.e., $\beta \rightarrow -\infty$. Otherwise, the optimal quantity remains uncertain. The ex post quantity adjustment also depends on the cost of public funds: the higher λ , the less sensitive is the total quantity to the cost shocks.

In turn, the allocation across technologies can be expressed as

$$q_1^{FB}(\theta_1, \theta_2) = \frac{Q^{FB}(\theta_1, \theta_2)}{2} + \frac{\Phi(\lambda, \Delta c + \Delta \theta)}{2}, \quad (3)$$

$$q_2^{FB}(\theta_1, \theta_2) = \frac{Q^{FB}(\theta_1, \theta_2)}{2} - \frac{\Phi(\lambda, \Delta c + \Delta \theta)}{2}, \quad (4)$$

where

$$\Phi(\lambda, \Delta c + \Delta \theta) = \frac{1}{\gamma} \frac{1 + \lambda}{1 + 2\lambda} (\Delta c + \Delta \theta) \quad (5)$$

captures the difference in the quantities allocated to the two technologies. This difference depends on the ex ante cost difference, Δc , as well as on the ex post cost shocks, $\Delta \theta$. The more efficient technology gets a higher allocation; and the more so the greater its cost advantage and the flatter the aggregate supply curve. However, whenever $\lambda > 0$, the quantity allocated to the more efficient technology is lower than under the cost-minimising solution (which would be obtained if $\lambda = 0$).²⁴ This departure is increasing in λ : the more concerned is the regulator about the cost of public funds, the more she is willing to distort her demand schedules away from the cost-efficient solution.

²³ As it will become clear, this case is the most favourable one for technology neutrality.

²⁴ Trivially, quantities would not be distorted if the ex post costs were identical, $c_1 + \theta_1 = c_2 + \theta_2$, which happens with probability zero.

As a consequence, even if the two technologies are perfectly symmetric on the benefit side, the prices for the two technologies differ whenever $\lambda > 0$. In particular, prices are given by

$$p_1^{FB}(\theta_1, \theta_2) = c_1 + \theta_1 + \frac{\gamma}{2}[Q^{FB}(\theta_1, \theta_2) + \Phi(\lambda, \Delta c + \Delta \theta)],$$

$$p_2^{FB}(\theta_1, \theta_2) = c_2 + \theta_2 + \frac{\gamma}{2}[Q^{FB}(\theta_1, \theta_2) - \Phi(\lambda, \Delta c + \Delta \theta)].$$

This finding should not come as a surprise as it reflects a standard third-degree price discrimination motive. Note that, similarly to quantities, prices adjust to cost shocks.

Adding asymmetries across technologies on the benefit side (i.e., letting $\partial B(q_t, q_{-t})/\partial q_t$ differ from $\partial B(q_t, q_{-t})/\partial q_{-t}$) would only change equilibrium prices and quantities through changes in demand schedules, as implicit in the term $\partial B(q_t, q_{-t})/\partial q_t$. This may result in more or less price divergence across the two technologies, and in a larger or smaller departure from the cost-minimising solution, depending on the regulator's preferences and the cost of supplying the different technologies. However, the key results would remain unchanged: the optimal mechanism departs from full cost efficiency and delivers two (technology-specific) prices, with or without differences on the benefit side. Furthermore, both prices and quantities would adjust to the cost shocks.

While the optimal auction has the great advantage of indirectly solving the regulator's information problem, in reality, at least in the realm of resource and renewable-energy auctions, it has rarely been used, if ever. For the most part, regulators tend to rely on simpler policy designs that adjust only partially to actual cost realisations, whether fixing quantities ex ante and letting prices adjust ex post or, alternatively, fixing prices ex ante and letting quantities adjust ex post. Some may argue that these simpler designs leave less room for ex post arbitrary adjustments. However, the optimal auction is also immune to such concerns as it commits the regulator to act upon a pre-announced schedule. It is arguable whether schedules are easier to 'manipulate' than quantities or prices, or the reverse.

Without delving into the political economy of why some instruments enjoy more support than others, in the rest of the paper we analyse procurement designs that have been used or proposed in practice, whether quantity or price based. In the presence of asymmetric information, none of these designs will approach the outcome of the optimal auction (Lemma 1); unless, of course, $\lambda = 0$ and $\eta = 1$, in which case the optimal auction converges to a technology-neutral auction. Therefore, our goal is to understand whether and under what conditions some instruments may be superior to others.

3. Quantity-Based Procurement

We start our analysis with the case in which the regulator chooses quantity targets. By construction, this approach departs from the optimal mechanism, given that it does not allow for ex post quantity adjustments. Does this approach differ from the optimal mechanism in other dimensions?

We first consider two quantity-based mechanisms, which are either technology neutral or technology specific. A technology-neutral auction is open to both technologies. Suppliers bid their true costs and the regulator pays them the market-clearing price times the total quantity. Provided the regulator is allowed to discriminate across bidders, an alternative is to run technology-specific auctions, with the regulator paying bidders according to two different market-clearing prices.

Which of these two approaches gets closer to the optimal mechanism? We still assume that both technologies are perfect substitutes, and leave the analysis of imperfect substitutes for the end of this section.

3.1 Technology-Neutral Auctions

Consider first a technology-neutral auction and denote by Q^N the regulator's optimal quantity choice:

$$Q^N = \arg \max_Q W(Q).$$

Given Q^N , the ex post allocation across technologies will depend on the realised cost shocks, even if the total quantity does not. Indeed, the quantity allocation will be such that the marginal costs of the two technologies will be equalised to the market-clearing price,

$$p^N = c_1 + \theta_1 + \gamma q_1^N = c_2 + \theta_2 + \gamma q_2^N. \quad (6)$$

Using (6), the equilibrium contribution of each technology to total output is given by

$$q_1^N(\theta_1, \theta_2) = \frac{Q^N}{2} + \frac{\Phi(0, \Delta c + \Delta \theta)}{2}, \quad (7)$$

$$q_2^N(\theta_1, \theta_2) = \frac{Q^N}{2} - \frac{\Phi(0, \Delta c + \Delta \theta)}{2}, \quad (8)$$

where the difference between the two quantities, $\Phi(0, \Delta c + \Delta \theta)$, is captured by (5) with $\lambda = 0$. Since the allocation across technologies fully adjusts to the ex ante cost difference and the ex post cost shocks, cost efficiency is achieved. This allocation is the same as under the optimal mechanism only when the regulator is not concerned about firms' rents.

Using (6) to (8), one can also obtain the market-clearing price as a function of the cost shocks,

$$p^N(\theta_1, \theta_2) = \frac{c_1 + c_2 + \theta_1 + \theta_2}{2} + \frac{\gamma}{2} Q^N.$$

The market-clearing price reaches the maximum level of uncertainty when shocks are perfectly and positively correlated, and the minimum when shocks are perfectly and negatively correlated (i.e., $\theta_1 = -\theta_2$), in which case, there is no price uncertainty.

In sum, technology neutrality departs from the optimal mechanism for two main reasons. First, the total quantity does not adjust to the cost shocks. And second, the quantity allocation across technologies achieves full cost efficiency at the cost of leaving too high rents. This is stated in our first lemma below.

LEMMA 2. Suppose that $\eta = 1$. In a technology-neutral auction,

(i) the total quantity is optimal ex ante, $Q^N = E[Q^{FB}]$, but not ex post,

$$Q^N = Q^{FB}(\theta_1, \theta_2) + \hat{Q}(\theta_1, \theta_2); \quad (9)$$

(ii) the quantities allocated to each technology are not optimal, neither ex ante nor ex post. In particular, it now becomes

$$\frac{E[\Delta q^{FB}]}{E[\Delta q^N]} = \frac{\Delta q^{FB}(\theta_1, \theta_2)}{\Delta q^N(\theta_1, \theta_2)} = \frac{1 + \lambda}{1 + 2\lambda} \leq 1.$$

Lemma 2 above points at two important results. First, under technology neutrality, the regulator procures the optimal total quantity in expected terms. The reason is that procuring an *extra unit*

of output is expected to cost the same to society as under the optimal mechanism, taking into account both the actual costs as well as the fiscal distortions. However, once the cost shocks are realised, the lack of flexibility of the quantity target gives rise to distortions as compared to the first best.

The second result of Lemma 2 shows that the technology allocation is sub-optimal; not only ex post, once the cost shocks are realised, but more interestingly, ex ante. In particular, as compared to technology neutrality, optimality calls the regulator to procure less of the efficient technology and more of the less efficient one to reduce payments.²⁵ Indeed, by increasing the allocation to the less efficient technology, the regulator can reduce the over-compensation to the more efficient technology. Since the reduction in the rents going to the efficient technology dominates over the increase in the rents going to the inefficient one, total payments decrease. Since technology neutrality fails to implement this outcome, it leads to inefficiently high social costs: the greater efficiency in the technology choice is more than offset by the social costs of the higher rents.

3.2 Technology-Specific Auctions

Consider now a mechanism that exploits the regulator's ability to discriminate suppliers according to their technologies. In particular, consider two technology-specific (uniform-price) auctions and denote by q_1^S and q_2^S the regulator's optimal choices, i.e.,

$$\{q_1^S, q_2^S\} = \arg \max_{q_1, q_2} W(q_1, q_2),$$

leading to $Q^S = q_1^S + q_2^S$. The market-clearing price in auction $t = 1, 2$, denoted p_t^S , is equal to the marginal cost of that technology,

$$p_t^S = c_t + \theta_t + \gamma q_t^S. \quad (10)$$

The regulator chooses the allocation in order to equate the expected marginal social costs across technologies,

$$(c_1 + \gamma q_1^S)(1 + \lambda) + \lambda \gamma q_1^S = (c_2 + \gamma q_2^S)(1 + \lambda) + \lambda \gamma q_2^S.$$

The expected marginal costs of the two technologies are equalised only when the regulator is not concerned about payments, i.e., $\lambda = 0$. Otherwise, the regulator takes into account the impact of the allocation on expected payments, as captured by the second term on both sides of the equality. Note that this expression does not depend on the realised cost shocks, θ_1 and θ_2 , as the regulator has to choose the technology targets ex ante.²⁶

Using $Q^S = q_1^S + q_2^S$, the equilibrium contribution of each technology to total output can be written as

$$q_1^S = \frac{Q^S}{2} + \frac{\Phi(\lambda, \Delta c)}{2}, \quad (11)$$

$$q_2^S = \frac{Q^S}{2} - \frac{\Phi(\lambda, \Delta c)}{2}, \quad (12)$$

²⁵ Note that if it were costless to raise public funds ($\lambda = 0$), there would be no quantity distortion under the optimal mechanism and the outcome would be the same as under technology neutrality.

²⁶ The regulator can improve upon this technology-specific design by running the two separate auctions sequentially. We leave for future work the study of which technology is better to auction first: the more uncertain (i.e., the one with higher σ) or the less efficient ex ante (i.e., the one with higher c). The choice will solve a trade-off between learning and ex post adjustment.

where the difference between the two quantities, $\Phi(\lambda, \Delta c)$, is captured by (5) with no adjustment to the realised cost difference, $\Delta\theta$.

The resulting prices are

$$p_1^S(\theta_1) = c_1 + \theta_1 + \frac{\gamma}{2}[Q^S + \Phi(\lambda, \Delta c)],$$

$$p_2^S(\theta_2) = c_2 + \theta_2 + \frac{\gamma}{2}[Q^S - \Phi(\lambda, \Delta c)].$$

Similarly to the optimal mechanism, the two prices need not coincide. However, unlike the optimal mechanism, the technology-specific prices depend exclusively on each technology's own cost shock. Our next lemma formally compares these two mechanisms.

LEMMA 3. Suppose that $\eta = 1$. In a technology-specific auction,

(i) the total quantity is optimal ex ante, $Q^S = E[Q^{FB}]$, but not ex post,

$$Q^S = Q^{FB}(\theta_1, \theta_2) + \hat{Q}(\theta_1, \theta_2); \quad (13)$$

(ii) the quantities allocated to each technology are optimal ex ante, but not ex post. In particular,

$$\frac{E[\Delta q^{FB}]}{E[\Delta q^S]} = 1 \quad \text{but} \quad \frac{\Delta q^{FB}(\theta_1, \theta_2)}{\Delta q^S(\theta_1, \theta_2)} = 1 + \frac{\Delta\theta}{\Delta c}.$$

It follows that the technology-specific approach departs from the optimal mechanism because quantities, both the total quantity as well as the technology-specific targets, do not adjust to the cost shocks. However, and unlike the technology-neutral approach, the ex ante allocation across technologies is optimal.

3.3 Technology-Neutral versus Technology-Specific Auctions

Having compared each auction format against the optimal mechanism, we now compare technology-neutral versus technology-specific auctions under the assumption of perfect substitutes. Lemmas 2 and 3 greatly facilitate the comparison. On the one hand, by looking at statements (i) in the lemmas, we can conclude that the total quantity procured is invariant to the auction format, $Q^N = Q^S$. The reason was alluded to already: since the expected marginal social costs are equalised at the margin, procuring an extra unit of output under either instrument is expected to cost the same to society.

On the other hand, by looking at statements (ii) in the lemmas, we can conclude that the quantities allocated to each technology differ across the two auction formats. Ex ante,

$$q_1^S - E[q_1^N] = E[q_2^N] - q_2^S = \frac{\Phi(\lambda, \Delta c) - \Phi(0, \Delta c)}{2} < 0,$$

quantities differ because the technology-specific approach, similarly to the optimal mechanism, allocates a greater share of total output to the less efficient technology in order to reduce rents. Ex post,

$$q_1^S - q_1^N(\theta_1, \theta_2) = q_2^N(\theta_1, \theta_2) - q_2^S = \frac{\Phi(\lambda, \Delta c) - \Phi(0, \Delta c + \Delta\theta)}{2},$$

quantities further differ because the allocation under the technology-specific approach does not adjust to cost shocks.

These two differences are key for the welfare analysis. Indeed, the comparison of payments and costs is fundamentally linked to these quantity distortions. Comparing expected payments,

$$E[T(q_1^S, q_2^S)] - E[T(q_1^N, q_2^N)] = \frac{\gamma}{2} [\Phi(\lambda, \Delta c) - \Phi(0, \Delta c)] \Phi(\lambda, \Delta c) < 0, \quad (14)$$

shows that payments are lower under the technology-specific approach, with the difference increasing in the expected quantity distortion, which in turn increases in λ . However, this reduction in expected payments comes at the expense of increasing expected costs, as captured by the first term of the right-hand side in the expression

$$E[C(q_1^S, q_2^S)] - E[C(q_1^N, q_2^N)] = \frac{\gamma}{4} [\Phi(\lambda, \Delta c) - \Phi(0, \Delta c)]^2 + \frac{E[(\Delta\theta)^2]}{4\gamma}. \quad (15)$$

The second term in (15) further captures the fact that, under cost uncertainty, costs are minimised under a technology-neutral approach as it allows for ex post adjustments.

Expressions (14) and (15) capture the basic rent-efficiency trade-off faced by the regulator who must decide whether to keep technologies competing together in the same auction or to separate them. The former approach favours cost efficiency while the latter allows for reducing payments. This trade-off is at the heart of our first proposition.

PROPOSITION 1. *Suppose that $\eta = 1$. The regulator should favour a technology-neutral auction over technology-specific auctions if and only if the difference in expected welfare is positive,*

$$W_q^N - W_q^S = \frac{1}{4\gamma} \left[E[(\Delta\theta)^2] - \frac{\lambda^2}{1 + 2\lambda} (\Delta c)^2 \right] > 0, \quad (16)$$

where $E[(\Delta\theta)^2] = 2(1 - \rho)\sigma_1\sigma_2 + (\sigma_1 - \sigma_2)^2$.

According to the proposition, the regulator should opt for the technology-neutral design when the expected efficiency loss of not doing so—as captured by the first term within brackets in (16)—is more important than the additional rents left with suppliers from not running separate auctions—as captured by the second term. Expression (16) tells us that a well-informed regulator (which here is equivalent to assuming that $\sigma_t \rightarrow 0$ for $t = 1, 2$, and hence $E[(\Delta\theta)^2] \rightarrow 0$) should always run separate auctions, with q_1^S and q_2^S chosen so as to balance the minimisation of costs and payments. A similar prescription should be followed if the two technologies are subject to similar shocks (i.e., $\rho \rightarrow 1$ and $\sigma_1 = \sigma_2$, again implying that $E[(\Delta\theta)^2] \rightarrow 0$), because in this case, ex post cost minimisation is no longer an issue.

As incomplete information mounts, however, she may reverse her decision in favour of technology neutrality unless the cost for the regulator of leaving rents with the suppliers is too large.²⁷ This ultimately depends on the amount of over-compensation to the more efficient suppliers—as captured here by the cost difference Δc —and the unit price of this over-compensation—as captured here by the shadow cost of public funds, λ (note that $\lambda^2/(1 + 2\lambda)$ is increasing in λ).

²⁷ This reversal is more likely, not only as ρ approaches -1 , but also as the level of uncertainty across the two technologies differs. To see the latter, suppose that $\sigma_1 + \sigma_2 = k$, where k is some constant. The first term within brackets in (16) reaches a maximum of k^2 , either when $\rho = -1$ and for any pair (σ_1, σ_2) , or when either σ_1 or σ_2 is equal to zero.

3.4 Imperfect Substitutes

Let us now consider the case of imperfect substitutes, $\eta \in [0, 1)$. Clearly, the solution under technology neutrality remains unchanged, given that the equilibrium market prices and quantities result from the equalisation of marginal costs across technologies, which do not depend on η . Matters become different under the optimal mechanism and under the technology-specific approach, given that in these cases the solution results from equalising the difference between marginal benefits and marginal costs across technologies. Since the former depend on η , the regulator's preferred technology allocation depends on η as well.

In particular, the difference between q_1^{FB} and q_2^{FB} is now given by

$$\Phi(\eta, \lambda, \Delta c + \Delta \theta) = \frac{1 + \lambda}{(1 + 2\lambda)\gamma - 2\beta(1 - \eta)} (\Delta c + \Delta \theta) \geq 0, \quad (17)$$

which generalises (5) by allowing for $\eta \leq 1$. Indeed, (17) boils down to (5) for the case of perfect substitutes, $\eta = 1$. Interestingly,

$$\frac{\partial \Phi(\eta, \lambda, \Delta c + \Delta \theta)}{\partial \lambda} < 0 < \frac{\partial \Phi(\eta, \lambda, \Delta c + \Delta \theta)}{\partial \eta}.$$

This means that moving away from the case of perfect substitutes, i.e., reducing η below 1, pushes the two technologies closer to each other, thus reinforcing the effect of allowing for costly public funds, i.e., increasing λ above zero. Intuitively, substituting one unit of the low-cost (and high-quantity) technology with one unit of the high-cost (and low-quantity) technology increases costs, but it also raises benefits relatively more, given the degree of imperfect substitution across technologies. The solution thus strikes a balance between reducing costs while simultaneously increasing social benefits.

Equivalently, the difference between q_1^S and q_2^S is now given by $\Phi(\eta, \lambda, \Delta c)$, which only differs from (17) in that the technology-specific targets are fixed ex ante and hence do not respond to the cost shocks $\Delta \theta$.

It is important to note that, while η affects the technology allocation, it does not affect the total quantity demanded. The reason is that the sum of the marginal benefits and marginal costs across technologies only depends on Q , regardless of how it is split among the two technologies. It follows that Q^{FB} and $Q^N = Q^S$ remain as in (2) and (9) or (13), respectively.

As a consequence, Lemma 3 comparing the technology-specific approach versus the optimal mechanism remains unchanged, regardless of η . In contrast, part (ii) of Lemma 2 comparing the allocation under technology neutrality versus the optimal mechanism changes when we allow for $\eta < 1$. In particular, we have

$$\frac{E[\Delta q^{FB}]}{E[\Delta q^N]} = \frac{\Delta q^{FB}(\theta_1, \theta_2)}{\Delta q^N(\theta_1, \theta_2)} = \frac{(1 + \lambda)\gamma}{(1 + 2\lambda)\gamma - 2\beta(1 - \eta)} < \frac{1 + \lambda}{1 + 2\lambda}.$$

In other words, the allocation under technology neutrality is not optimal because it minimises costs without taking into account the impact on social benefits. This effect adds to the distortion created by the rent-efficiency trade-off highlighted before.

The above results have implications for the comparison between the technology-neutral and the technology-specific approaches, as it is fundamentally linked to quantity distortions. Indeed, following the same approach as before, the difference in payments and costs across the two approaches can be expressed as a function of the Φ functions. Noting that the allocation under technology neutrality is the same as if public funds were costless ($\lambda = 0$) and technologies were

perfect substitutes ($\eta = 1$) allows us to write the difference in payments as

$$E[T(q_1^S, q_2^S)] - E[T(q_1^N, q_2^N)] = \frac{\gamma}{2} [\Phi(\eta, \lambda, \Delta c) - \Phi(1, 0, \Delta c)] \Phi(\eta, \lambda, \Delta c) < 0. \quad (18)$$

Since $\Phi(\eta, \lambda, \Delta c)$ is increasing in η , the difference in payments grows larger as η falls below 1.

For the same reason, the difference in costs is enlarged when technologies are imperfect substitutes,

$$E[C(q_1^S, q_2^S)] - E[C(q_1^N, q_2^N)] = \frac{\gamma}{4} [\Phi(1, 0, \Delta c) - \Phi(\eta, \lambda, \Delta c)]^2 + \frac{E[(\Delta\theta)^2]}{4\gamma} > 0. \quad (19)$$

Adding to this trade-off is the impact on the social benefit side, which is no longer the same under the two approaches,

$$\begin{aligned} E[B(q_1^S, q_2^S)] - E[B(q_1^N, q_2^N)] \\ = -(1 - \eta) \frac{\beta}{2} [\Phi^2(1, 0, \Delta c + \Delta\theta) - \Phi^2(\eta, \lambda, \Delta c)] \\ > 0. \end{aligned} \quad (20)$$

This means that, unless technologies are seen as perfect substitutes ($\eta = 1$), benefit considerations tilt the regulator's trade-off. In particular, since the term in square brackets is positive and $\beta < 0$, when technologies are seen as imperfect substitutes ($\eta < 1$), the regulator should lean toward technology separation.

Our next proposition, which generalises Proposition 1 for the case of perfect and imperfect substitutes, clarifies how this trade-off is ultimately resolved in the general case with costly public funds and (possibly) imperfect substitutes.

PROPOSITION 2. *Suppose that $\eta \in [0, 1]$. The regulator should favour a technology-neutral auction over technology-specific auctions if and only if the difference in expected welfare is positive,*

$$W_q^N - W_q^S = \frac{1}{4\gamma} [\Upsilon(\eta)E[(\Delta\theta)^2] - \Psi(\eta)(\Delta c)^2] > 0,$$

where $\Upsilon(\eta) = 1 + 2\beta(1 - \eta)/\gamma$ and

$$\Psi(\eta) = \frac{1}{\gamma} \frac{(\lambda\gamma - 2\beta(1 - \eta))^2}{(1 + 2\lambda)\gamma - 2\beta(1 - \eta)} > 0.$$

Expression $\Psi(\eta)$ is decreasing in η and equals $\lambda^2/(1 + 2\lambda)$ when $\eta = 1$, the case of perfect substitutes (as in Proposition 1). And since $\Upsilon(\eta)$ is increasing in η , allowing for imperfect substitutes unambiguously favours separation. However, this does not necessarily mean that, as η drops, separation eventually dominates technology neutrality, regardless of the degree of asymmetric information. Indeed, even when technologies are completely unrelated on the benefit side ($\eta = 0$), technology neutrality can still be the preferred choice when the benefit curves are rather flat (i.e., β not too negative). The reason is that the marginal benefit of both technologies would then be roughly the same, bringing us back to Proposition 1: the cheapest solution should then be preferred, and this sometimes involves technology neutrality. Obviously, introducing asymmetries across technologies on the benefit side would further favour separation.

4. Hybrid Schemes

Since neither technology neutrality nor technology separation achieves optimality, one may argue in favour of hybrid approaches that allow for some partial separation between technologies. We consider two approaches currently in use: technology banding and MTQs. To simplify the exposition, and without much loss in insights, here and in the following sections we assume that $\sigma_1 = \sigma_2 = \sigma$ and $\eta = 1$.

4.1 Technology Banding

A handful of countries currently rely on technology banding for setting renewable support.²⁸ The idea is to run a single uniform-price auction with suppliers of the ex ante inefficient technology (or less resourceful location) receiving a handicap to compete more effectively with suppliers of the ex ante more efficient technology or location.

Let $\alpha > 1$ be the handicap received by the ex ante inefficient technology (technology 2). This means that if p^B is the market-clearing price under banding, technology 2 gets a price of αp^B for each unit supplied, while technology 1 just gets p^B . Thus, at every price, suppliers of technology 2 are willing to offer a greater quantity the higher the handicap α .²⁹

The regulator's optimal banding choice is

$$\{\alpha^B, Q^B\} = \arg \max_{\alpha, Q} W(\alpha, Q),$$

where $Q^B = q_1^B + q_2^B$. From the market clearing condition,

$$p^B = c_1 + \theta_1 + \gamma q_1^B = \frac{1}{\alpha^B} (c_2 + \theta_2 + \gamma q_2^B),$$

one can obtain the equilibrium contribution of each technology,

$$q_1^B(\alpha^B; \theta_1, \theta_2) = \frac{Q^B}{1 + \alpha^B} + \frac{c_2 + \theta_2 - \alpha^B(c_1 + \theta_1)}{(1 + \alpha^B)\gamma}, \quad (21)$$

$$q_2^B(\alpha^B; \theta_1, \theta_2) = \frac{\alpha^B Q^B}{1 + \alpha^B} - \frac{c_2 + \theta_2 - \alpha^B(c_1 + \theta_1)}{(1 + \alpha^B)\gamma}. \quad (22)$$

In turn, the equilibrium market-clearing price as a function of the shocks is given by

$$p^B(\alpha^B; \theta_1, \theta_2) = \frac{c_1 + c_2 + \theta_1 + \theta_2}{1 + \alpha^B} + \frac{\gamma}{1 + \alpha^B} Q^B. \quad (23)$$

Since α can always be set equal to one (and Q^B equal to Q^N), the banding design is by construction (weakly) superior to the technology-neutral design. The only case when the two

²⁸ One example of technology banding is provided by the reference yield model for wind that has been in place in Germany since 2000. It relies on plant- and site-specific adjustment factors that favour investment in sites with less wind. The Renewable Obligation scheme that was in place in the United Kingdom (which was very similar to the Renewable Portfolio Standard programs in the United States) offers another example. Renewable producers are allowed to issue Renewable Obligation Certificates (ROCs) that electricity suppliers have to buy to meet their obligations. While the default was that one ROC would be issued for each MWh of renewable output, the system was subsequently reformed so that some technologies were allowed to issue more, others less. For instance, in 2017, installations were entitled to receive 1.8 ROCs per MWh of offshore wind, 0.9 ROCs for onshore wind installations and 1.4 ROCs for building-mounted solar photovoltaics (UK Government, 2013).

²⁹ This price adjustment is also often used whenever the two goods are considered to be of different qualities, e.g., liquidity auctions, backed by strong or weak collateral. In this case, the high-quality good is given a handicap or a supplement on top of the market price. In the product-mix auction (Klemperer, 2010), the handicap is endogenously determined, together with the fraction of high-quality goods, according to the regulator's demand schedule.

designs converge is when $\Delta c = 0$. Less evident is whether a banding design can also be superior to a technology-specific design, and if so, under what circumstances. To explore this possibility, it helps to start with the following intermediate result.

LEMMA 4. *In the absence of uncertainty, i.e., $\sigma \rightarrow 0$, (i) the banding design replicates the technology-specific design, with $Q^B = q_1^S + q_2^S$ and $\alpha^B = p_2^S/p_1^S$, and (ii) either design strictly dominates technology neutrality, i.e., $W_q^B = W_q^S > W_q^N$.*

PROOF. It follows immediately from comparing W_q^B and W_q^S when $\theta_1 = \theta_2 = 0$ and from Proposition 1. $\square$

In the absence of uncertainty, the regulator is indifferent between technology banding and technology separation since in either case she has two instruments at her disposal. Matters change, however, as we introduce uncertainty. One may speculate that, under uncertainty, one should lean in favour of the banding option since, by allowing for some technology substitution, it appears better equipped at containing total costs. But, akin to Proposition 1, allowing for this substitution may come at the expense of leaving higher rents with suppliers, to the extent that technology separation may nevertheless prevail as the best option.

PROPOSITION 3. *Suppose that technology-specific auctions are superior to technology-neutral auctions, i.e., $W_q^S > W_q^N$. There exists a correlation cutoff, $\bar{\rho} < 1$, above which technology-specific auctions also dominate technology banding, i.e., $W_q^S > W_q^B$.*

PROOF. See Appendix B. $\square$

To convey the intuition of Proposition 3, let us go through some key steps of the proof. To start, note that there is no point in comparing technology banding to technology separation if the latter is dominated by technology neutrality. In that case, banding would be automatically superior, by construction. Therefore, suppose that λ is large enough so that technology separation dominates technology neutrality, i.e., (16) in Proposition 1 does not hold.

Building from Lemma 4, suppose for now that $Q^B = q_1^S + q_2^S$ for any level of uncertainty (we will shortly comment on this). This reduces the comparison between banding and separation to one dimension: how uncertainty affects expected costs and payments across designs. Under technology separation, expected costs and payments are invariant to uncertainty (see Subsection 3.2). Hence, we just need to understand how uncertainty affects expected costs and payments under banding. Assuming that $\alpha^B = E[p_2^S]/E[p_1^S]$, we can use (21) and (22) to obtain expressions for these two components as

$$E[C^B(Q^B, \alpha^B)] = E[C^S(q_1^S, q_2^S)] + \frac{\sigma^2[\rho(1 + (\alpha^B)^2) - 2\alpha^B]}{\gamma(1 + \alpha^B)^2} \quad (24)$$

and

$$E[T^B(Q^B, \alpha^B)] = E[T^S(q_1^S, q_2^S)] + \frac{\sigma^2(1 + \rho)(\alpha^B - 1)^2}{\gamma(1 + \alpha^B)^2}, \quad (25)$$

where $Q^B = q_1^S + q_2^S$. Consistent with Lemma 4, as $\sigma \rightarrow 0$ (and $\alpha^B \rightarrow p_2^S/p_1^S$), expected costs and payments converge across the two formats so that they become no different.

As we increase σ , however, two things occur: expected costs can go up or down, depending on ρ and α^B , and expected payments can only go up, except when $\rho = -1$. To be more precise

about the implications for the welfare comparison, it helps to focus on two extreme values of ρ . Consider first the case of perfectly and negatively correlated cost shocks, i.e., $\rho = -1$. From (24) and (25), banding is unambiguously superior to separation because expected costs are lower under banding, while expected payments are the same as under separation. It is easy to understand why payments coincide: when $\rho = -1$, the market-clearing price under banding (23) becomes certain (just like the market-clearing price under separation), thereby making the regulator's expected payments certain as well.

On the other hand, expected costs are lower under banding because it allows for substitution across technologies, albeit incompletely since $\alpha^B > 1$ when it is most valuable from a cost containment point of view. Interestingly, the value of this substitution is complete at $\rho = -1$, despite $\alpha^B > 1$. In fact, expected cost savings under banding relative to separation, which add to σ^2/γ , are exactly the same as under technology neutrality relative to separation (see Proposition 1). However, as ρ departs from -1 , cost savings under banding are not as large as under technology neutrality because of the efficiency distortion introduced by setting $\alpha^B > 1$.

Consider now the other correlation extreme, $\rho = 1$. Unlike the previous case, separation is now unambiguously superior to banding because both expected costs, as well as expected payments, are lower under separation. The fact that payments are higher under banding is not very surprising because $\rho = 1$ gives rise to highly uncertain market-clearing prices, leading to highly uncertain payments. More intriguing is the fact that banding fails to provide any cost containment at all. Part of the reason for this was already alluded to in the previous paragraph. From Proposition 1, we know that allowing for technology substitution when $\rho = 1$ does not provide any cost containment advantage at all. The problem with banding, however, is that technology substitution is distorted by the fact that $\alpha^B > 1$. And this distortion has a price. From (21) and (22) we can see that, under a positive cost shock, $\theta_1 = \theta_2 > 0$, quantities procured of each technology move further away from their cost-minimising levels (q_2 moves further up and q_1 further down). Under a similar, but negative cost shock, quantities move instead closer to their cost-minimising levels. But costs are convex, so the first effect dominates the second, as Jensen's inequality predicts. If α^B were equal to one, these two effects would cancel each other out.³⁰

Going over these extreme correlation scenarios allows us to establish, by continuity, the existence of a correlation cutoff $\bar{\rho} < 1$ that leaves the regulator indifferent between technology separation and banding. Using the regulator's indifference condition, $W_q^S = W_q^B$, this cutoff is given by³¹

$$\bar{\rho} = \frac{2\alpha^B - \lambda(\alpha^B - 1)^2}{1 + (\alpha^B)^2 + \lambda(\alpha^B - 1)^2} < 1. \quad (26)$$

For $\rho > \bar{\rho}$, separation dominates banding, and vice versa.

One key component in the cutoff expression (26) is the cost of public funds, λ . A lower value of λ pushes $\bar{\rho}$ further up, making banding more attractive. The reason is that the regulator's

³⁰ As we explain in Appendix B, the case of $\rho = 1$ requires an additional step before one can formally establish that $W_q^S > W_q^B$. Unlike when $\rho = -1$, both α^B and Q^B are indeed not invariant to the introduction of uncertainty, which implies that the deterministic component in W_q^B is no longer equal to $W_q^S = W(q_1^S, q_2^S)$. But since, under separation, q_1 and q_2 can always be chosen to exactly replicate the deterministic component in W_q^B , it must be true that the deterministic component in W_q^B falls with uncertainty. Hence, the superiority of separation at $\rho = 1$ is only reinforced as we introduce uncertainty.

³¹ Note that this cutoff expression is strictly valid as $\sigma \rightarrow 0$. As σ increases, two things happen: α^B goes down and the deterministic component of W_q^B also goes down. These factors act in opposing directions, but in Appendix B we show that the first factor dominates, so $\bar{\rho}$ goes up with uncertainty, but remains away from 1.

payments do not weigh as much, thereby mitigating the advantage of separation in reducing rents. The other key component in (26) is α^B . A lower value of α^B also pushes $\bar{p}$ further up, making banding more attractive. Again, a lower α^B means that rent extraction is less important and that the potential cost distortions from imperfect substitution across technologies under banding will not be as large.

The factors that contribute to a lower α^B are very intuitive as well. As shown in Appendix B, α^B is weakly decreasing with uncertainty, which is when (cost) efficiency considerations become more important, thereby enhancing the value of banding. In the same Appendix we also show that, as uncertainty vanishes, α^B reduces to³²

$$\alpha^B(\sigma \rightarrow 0) = 1 + \frac{2\lambda\Delta c}{\Delta c(1 + \lambda) + \gamma Q^B(1 + 2\lambda)} < \frac{5}{3},$$

which serves to show that α^B falls with lower values of λ and Δc and higher values of γ . Lower values of λ and Δc make rent extraction less important, the former by lowering its weight in the regulator's problem, the latter by reducing its magnitude. Last, a high γ also favours a lower α^B because the cost distortions are far costlier under a more convex cost curve.

4.2 Minimum Technology Quotas

Instead of relying on technology banding, Spain has introduced MTQs into its latest renewable auction. In our setting, an MTQ auction is a single uniform-price auction that ensures that each technology gets a minimum quota. When these MTQs are not binding, the auction reduces to a standard technology-neutral auction, with all technologies receiving the same price. As soon as one of the MTQs is binding, the binding technology receives a higher price as compared to that of the other technology.

Let $\underline{q}_t$ be the MTQ for technology $t = 1, 2$ and Q the total number of units to be auctioned off, with $\underline{q}_1 + \underline{q}_2 \leq Q$. When $\underline{q}_t$ is binding, $q_t = \underline{q}_t$ and $q_{-t} = Q - \underline{q}_t$, leading to a price wedge, $p_t = c_t + \theta_t + \gamma \underline{q}_t > p_{-t} = c_{-t} + \theta_{-t} + \gamma(Q - \underline{q}_t)$. Unlike technology banding, an MTQ can replicate the outcome of technology-neutral auctions, by setting $\underline{q}_1 = \underline{q}_2 = 0$ and $Q = Q^N$, and technology-specific auctions, by setting $\underline{q}_1 = q_1^S$, $\underline{q}_2 = q_2^S$ and $Q = \underline{q}_1 + \underline{q}_2 = Q^S$. Since technology banding fails to replicate the latter (see Proposition 3), one may be tempted to conclude that an MTQ is always superior to technology banding. We next show this is not necessarily the case.

For a given MTQ design, i.e., a triplet $\{\underline{q}_1^M, \underline{q}_2^M, Q^M\}$, the outcome may fall into three different regions, depending on the realisations of θ_1 and θ_2 : (i) the region where $\underline{q}_1$ is binding, that is, when $C'_1(\underline{q}_1; \theta_1) \geq C'_2(Q - \underline{q}_1; \theta_2)$ or

$$\theta_1 - \theta_2 \geq \Delta c + \gamma(Q - 2\underline{q}_1) \equiv \ell_1,$$

(ii) the region where $\underline{q}_2$ is binding, that is, when $C'_2(\underline{q}_2; \theta_1) \geq C'_1(Q - \underline{q}_2; \theta_1)$ or

$$\theta_1 - \theta_2 \leq \Delta c + \gamma(2\underline{q}_2 - Q) \equiv \ell_2,$$

and (iii) the neutrality region, that is, when

$$\ell_2 \leq \theta_1 - \theta_2 \leq \ell_1.$$

³² Note from (8), for example, that an interior solution—such that both technologies are always procured in equilibrium—requires $\gamma Q^B > \Delta c$, setting an upper bound for α^B of $5/3$.

Therefore, the optimal MTQ design can be found as the solution of the problem

$$\begin{aligned} \max_{\underline{q}_1, \underline{q}_2, Q} \sum_{t=1,2} \int_{\underline{\theta}_{-t}}^{\bar{\theta}_{-t}} \int_{\ell_t + \theta_{-t}}^{\bar{\theta}_t} W(\underline{q}_t, Q - \underline{q}_t; \theta_t, \theta_{-t}) g(\theta_t, \theta_{-t}) d\theta_t d\theta_{-t} \\ + \int_{\underline{\theta}_2}^{\bar{\theta}_2} \int_{\ell_2 + \theta_2}^{\ell_1 + \theta_2} W(q_1^N(Q, \theta_1, \theta_2), q_2^N(Q, \theta_1, \theta_2); \theta_1, \theta_2) g(\theta_1, \theta_2) d\theta_1 d\theta_2, \end{aligned}$$

where $W(\cdot)$ is the relevant welfare function for each region, and $q_1^N(Q, \theta_1, \theta_2)$ and $q_2^N(Q, \theta_1, \theta_2)$ are given by the quantity expressions under technology neutrality, (7) and (8), respectively.

Without solving this maximisation problem, it is not difficult to see that the optimal MTQ design may involve a single region, i.e., the region where the MTQ for the ex ante less efficient technology is binding (region (ii) in our case), or the three regions described above. The neutrality region only exists as a transition between regions (i) and (ii), i.e., when shocks are sufficiently large relative to Δc so that one technology becomes more efficient than the other for some realisations of θ_1 and θ_2 , but not for others. This insight leads to the following result.

PROPOSITION 4. *Technology banding can be superior to MTQs and vice versa.*

PROOF. Using an example, suppose that cost shocks are such that $\Delta c > \bar{\theta}_1 - \underline{\theta}_2$. Since technology 1 is more efficient than technology 2 for any realisation of θ_1 and θ_2 , the optimal MTQ design reduces to technology-specific auctions; it only includes region (ii). And we know from Proposition 3 that in this case there exists a correlation cutoff, $\bar{\rho} < 1$, below which technology banding dominates technology-specific auctions. $\square$

Proposition 4 serves to illustrate that technology banding may be the right choice when one technology is clearly more efficient than the other, both from an ex ante as well as from an ex post perspective. In contrast, MTQ is more flexible in handling the regulator's uncertainty when it is hard to tell ex ante which technology will end up being more efficient, i.e., when cost shocks are large relative to Δc . In this case, MTQs are better at handling very different outcomes; namely, the fact that one technology may be more efficient than the other (regions (i) and (ii)), or that the two technologies may turn out to be equally efficient (the neutrality region (iii)).

5. Price-Based Procurement

So far, we have considered a regulator who procures a total of Q units of some good under different auction formats. While we have worked under the assumption that Q is chosen to maximise the welfare expression (1), all our results go through if Q is not under the regulator's control, but rather exogenously given. The case of an endogenous Q opens a new set of questions, however. In particular, it may no longer be preferable to rely on the quantity-based instruments we have considered so far, but rather on price-based instruments. In the presence of uncertainty, this gives rise to a new trade-off: under a quantity-based instrument, the total quantity is fixed, but prices adjust to shocks, whereas under a price-based instrument, prices are fixed, but quantities adjust to shocks. Recall that, under the optimal mechanism, both the total quantity as well as the prices adjust ex post.

If the regulator cannot discriminate across the different technologies, the best she can do within the family of price-based instruments is to post a single price at which she is ready to buy

whatever is supplied by each technology. But if she can discriminate suppliers according to their technologies, as assumed throughout, she can do better by posting two prices, p_1 and p_2 .

Since two prices are, by construction, superior to a single price (unless $\lambda = 0$, in which case they are welfare equivalent), the regulator's optimal pricing choice is

$$\{p_1^*, p_2^*\} = \arg \max_{p_1, p_2} W(q_1(p_1), q_2(p_2)),$$

where quantities $q_t(p_t, \theta_t)$ adjust so that prices equal marginal costs

$$p_t^* = c_t + \theta_t + \gamma q_t$$

for $t = 1, 2$. In expected terms, this price is analogous to (10), $p_t^* = E[p_t^S] \equiv c_t + \gamma q_t$, confirming that, under certainty, a regime of two separate prices is not different from a regime of two separate quantities.

The welfare comparison between prices and quantities yields the following proposition.

PROPOSITION 5. *Two posted prices dominate two technology-specific auctions if and only if*

$$W_p^S - W_q^S = \frac{\sigma^2(1 + \rho)}{\gamma^2} \left(\beta + \frac{\gamma}{2} \frac{2}{1 + \rho} \right) > 0. \quad (27)$$

PROOF. See Appendix B. □

When shocks are perfectly correlated, $\rho = 1$, (27) reduces to nothing, but Weitzman's (1974) seminal 'prices versus quantities' expression (just note that $\gamma/2$ is the combined slope of two supply curves, each with slope γ). The intuition of his result is well known: a relatively more convex supply curve favours prices because 'mistakes' on the supply side are costlier than on the benefit side. This analogy with Weitzman (1974) should not be surprising, as $\sigma_1 = \sigma_2 = \sigma$, $\rho = 1$ and $\eta = 1$ imply that the two technologies behave just as one.

As we move away from the perfect correlation case, however, the price instrument performs better than the quantity instrument, i.e., the difference $W_p^S - W_q^S$ is more likely to be positive (recall that $\beta < 0$). For imperfectly correlated shocks, prices allow the quantities allocated to the various technologies to better adjust ex post to the cost shocks, which helps to contain production costs while reducing uncertainty on the benefit side. Thus, because of technology substitution, the slope of the relevant marginal cost curve becomes flatter under price regulation, thereby favouring the price approach. In fact, when shocks exhibit similar variances and are perfectly and negatively correlated, $\rho \rightarrow -1$, prices are unambiguously superior to quantities because there is no longer uncertainty on the benefit side.³³

With two prices or two quantities, expected government payments are independent of the degree of cost correlation ρ and uncertainty σ . Since, under certainty, prices and quantities are equally suited to reduce suppliers' rents, it follows that, under uncertainty, expected government payments are also the same with two prices or two quantities, which explains why λ is absent from (27). This result does not mean, however, that price regulation should always be preferred to quantity regulation when (27) holds. It may still be optimal to opt for quantity regulation, in particular, for a technology-neutral auction. According to our next proposition, this may happen when λ is relatively small.

³³ While this multiple-technology analysis was already in Weitzman (1974), it is unclear why he compares technology-specific prices and quantities, given that in the absence of costly public funds technology separation brings no additional benefit. When $\lambda = 0$, technology neutrality dominates separation, strictly so under quantity regulation (Proposition 1) and weakly so under price regulation. Hence, the only meaningful comparison is between a single quantity and a single price.

PROPOSITION 6. *Two posted prices dominate a technology-neutral auction if and only if*

$$W_p^S - W_q^N = \frac{\lambda^2}{1 + 2\lambda} \left(\frac{\Delta c}{2\gamma} \right)^2 + \frac{\sigma^2(1 + \rho)}{\gamma^2} \left(\beta + \frac{\gamma}{2} \right) > 0. \quad (28)$$

PROOF. Immediate from the proofs of Propositions 1 and 5. $\square$

To convey some intuition, it helps to decompose $(W_p^S - W_q^N)$ in two terms: $(W_p^S - W_p^N) + (W_p^N - W_q^N)$. The first term, $(W_p^S - W_p^N)$, is the rent-extraction gain from using two prices as opposed to a single price. This is exactly captured by the first term in (28). The second term, $(W_p^N - W_q^N)$, is the Weitzman's trade-off between using a (single) price and a (single) quantity. This is exactly captured by the second term in (28).

Since we know from Proposition 1 that in the absence of costly public funds, a single quantity accommodates better to shocks than two quantities, $W_q^N > W_q^S$, it is clear that in such a case $W_p^S > W_q^N$ implies that $W_p^S > W_q^S$. With costly public funds, however, $W_p^S > W_q^N$ no longer implies that $W_p^S > W_q^S$. Indeed, when λ is not too large (meaning that the main objective is to minimise costs), it can well be the case that a technology-neutral auction dominates over the rest, $W_q^N > W_p^S > W_q^S$. The reason is that, while two prices allow for more quantity adjustment than two quantities, technology neutrality is the only instrument that allows quantities to fully adjust.

6. Application: Spain's 2021 Renewable Auction

Motivated by Spain's 2021 renewable auction, in this section we illustrate the use of our theory for policy analysis. Our application is not intended to provide an ex post empirical evaluation of the Spanish auction, since much of the data required to carry out such an exercise (e.g., actual bids and projects' locations) have not been made publicly available yet. Instead, we put ourselves in the following situation. Suppose that we were standing in 2020, and were asked about the pros and cons of auctioning 3,000 MW of renewable energy under an MTQ approach, as Spain actually did in January 2021.

To answer this question, we have collected the most representative data available at the time: detailed information on solar and wind investments undertaken in Spain during 2019. With these data, we look at the implications for social costs (investment efficiency and payments) of auctioning 3,000 MW under either: (i) the optimal mechanism, (ii) a single technology-neutral auction, (iii) two separate technology-specific auctions, (iv) a technology-neutral auction combined with technology banding or (v) a technology-neutral auction combined with MTQs. Since (iv) dominates (ii) and (v) dominates both (ii) and (iii), we are ultimately interested in the comparison between (iv) and (v) relative to the optimal mechanism (i).³⁴

Before moving on to the application, it is important to note that we will be departing from our theory model in two respects. First, the aggregate marginal cost function of each technology will not be necessarily linear, but most likely a step-wise function capturing the costs and sizes of the various investment projects (below we describe how we estimate these costs). And second, for the previous reason, the slopes of the marginal cost functions will not be constrained to be the same for the two technologies.

³⁴ Note that if the regulator were not constrained to procure a fixed quantity, the resulting departure from the optimal mechanism would be greater than that reported below.

6.1 Estimating supply curves

Our data set—all renewable investment projects undertaken in Spain during 2019—specifies several project characteristics, namely, their technology (either solar PV or wind), their maximum production capacity and their location, among others.³⁵ Using historic data on renewable production across the 50 Spanish provinces,³⁶ we have computed the expected production of each investment project over its lifetime (which we assume is equal to 25 years).³⁷ We denote it as q_{itl} for project i of technology t located in province l . A project's (long-run) average cost is given by the ratio between its investment cost and its expected production. By ranking projects of the same technology in increasing average-cost order, we construct the aggregate (long-run) supply curve of such technology.

We parametrise the investment cost of each project as $(c_t + \zeta \theta_t) k_i^\xi$, where c_t is the cost parameter of technology t , θ_t is a cost shock for technology t and k_i is the capacity of project i .³⁸ We set ξ equal to 0.9 to capture mild scale economies.³⁹ Regarding the parameter c_t , we set it up so that the average costs of all the projects in our sample equal the average costs of that technology, as reported by the International Renewable Association (IRENA) for 2018.⁴⁰ Even if average costs are set at this level, heterogeneity in locations and plant sizes gives rise to variation in average costs across projects.

Regarding the cost shock θ_t , we assume that it is distributed according to a standard normal distribution, with a correlation coefficient ρ across the cost shocks for the two technologies. To understand the role of cost correlation, we use three alternative assumptions: $\rho \in \{-0.8, 0, 0.8\}$. The parameter ζ simply allows us to change the weight of cost shocks on total costs; we set it equal to 1,800.⁴¹ For each value of ρ , we consider 100 independent draws of the cost parameters (θ_1, θ_2) , i.e., for solar and wind. For comparability purposes, we use the same realisations for all auction designs. Following Laffont and Tirole (1993, p. 38) and Laffont (2005, p. 15), we also allow for three possible values for the cost of public funds: $\lambda \in \{0, 0.2, 0.4\}$.

Figure 1 plots the expected supply curve, i.e., for the pair (0, 0) of cost shocks. As can be seen, the average costs of solar plants (denoted by circles) tend to be lower than the average costs of wind plants (denoted by triangles). However, the average cost curve of solar plants becomes very steep as we approach the capacity constraint, given that the most expensive projects are the small ones located in the least sunny regions. The average cost curve of wind plants tends to be higher, but flatter, as all wind projects tend to be similar in size and they tend to be located in the windiest

³⁵ Data source: Registry of Renewable Installations in Spain (RIPRE), made available by the Spanish Ministry of Environment (MITECO, 2022).

³⁶ These data are obtained from Red Eléctrica de España (2022), which is the Spanish electricity system operator.

³⁷ If instead, we assume a shorter lifetime, say, of 20 years, the main conclusions of this analysis would remain unchanged as long as we apply that number to both technologies.

³⁸ Note that in the model described in Section 1 we had implicitly assumed that all projects had unit capacity, $k_i = 1$. This difference is inconsequential, but allows us to introduce scale economies in project size.

³⁹ Setting $\xi = 1$ would imply that differences in the average cost of each project would only arise due to their different locations. Setting ξ at lower values would make the average cost curves steeper, while the average cost would remain fixed at the same value reported by IRENA (2020).

⁴⁰ In detail, IRENA reports that the investment cost of solar PV was 1,113 \$/kW and 1,833 \$/kW for wind (we use an exchange rate \$/euro equal to 1.12). These parameters come from IRENA's (2019) report for Germany (no cost is reported for the investment cost of solar PV in Spain).

⁴¹ In Appendix D, Table D.1, we report the results when the regulator faces smaller uncertainty about the costs. In particular, we set $\zeta = 900$. It can be seen that all our results regarding the ranking across formats remain unchanged. Furthermore, the comparison between Tables 1 and D.1 shows that all four mechanisms perform worse when there is more uncertainty, as expected.




Fig. 1. *The Expected Aggregate Supply Curve for Solar and Wind Investment Projects in Spain, 2019.*

Notes: This figure displays the expected (long-run) average cost curves for solar (circles) and wind (triangles) projects.

regions only. Note also that, according to the figure, it would be cost effective to procure 3,000 MW from both solar and wind projects.

6.2 Results

Table 1 summarises the results (expected social costs, including expected cost and payments) relative to the optimal mechanism, for nine (ρ, λ) pairs.

As shown in the table, technology neutrality gives rise to lower expected costs as compared to the optimal mechanism, at the cost of increasing payments. The resulting social costs are thus higher (between 6% and 20%, depending on ρ and λ). Technology neutrality performs relatively worse when the cost correlation is negative and the cost of public funds is high. It is optimal only when $\lambda = 0$, as expected.

The technology-specific approach reduces payments relative to the optimal mechanism at the cost of increasing costs, i.e., the technology-specific approach results in too much separation across technologies, as the allocation does not adjust to the cost shocks. As a result, the departure in social costs relative to the optimal mechanism can reach 17% when the regulator does not care about firms' rents. As expected, the social costs are closer to optimal when the cost correlation is positive, as the gains from adjusting quantities ex post are relatively small.

As compared to technology neutrality, costs under the technology-specific approach are always higher, while payments are lower. On the one hand, the relative cost inefficiency of technology separation increases for higher values of λ , as the quantity distortion gets larger. On the other, this also enlarges the payment gap between the two approaches. Overall, this trade-off tends to favour the technology-specific over the technology-neutral approach as it most often gives rise to lower social costs. However, there are four exceptions to this result: the three cases with $\lambda = 0$ and the case with $\lambda = 0.2$ and $\rho = -0.8$. In line with Proposition 1, this shows that one may

Table 1. *Simulation Results Relative to the Optimal Mechanism—Greater Uncertainty in Costs.*

ρ	λ	Social Costs				Costs				Payments			
		Neutral	Specific	Banding	MTQ	Neutral	Specific	Banding	MTQ	Neutral	Specific	Banding	MTQ
-0.8	0	1.0000	1.1730	1.0000	1.0000	1.0000	1.1730	1.0000	1.0000	1.3279	0.8670	1.3279	1.3279
-0.8	0.2	1.1198	1.1541	1.1019	1.0202	0.9941	1.1660	1.0034	0.9978	1.6855	1.1005	1.5454	1.1210
-0.8	0.4	1.2080	1.1471	1.1713	1.0368	0.9910	1.1624	1.0060	0.9948	1.7031	1.1120	1.5485	1.1327
0	0	1.0000	1.1030	1.0000	1.0000	1.0000	1.1030	1.0000	1.0000	1.1832	0.8000	1.1832	1.1832
0	0.2	1.0914	1.0850	1.0728	1.0136	0.9933	1.0956	1.0047	0.9960	1.5344	1.0375	1.3804	1.0931
0	0.4	1.1604	1.0787	1.1210	1.0213	0.9905	1.0925	1.0019	1.0069	1.5487	1.0471	1.3932	1.0540
0.8	0	1.0000	1.0123	1.0000	1.0000	1.0000	1.0123	1.0000	1.0000	1.0536	0.8386	1.0536	1.0536
0.8	0.2	1.0604	1.0104	1.0398	1.0050	0.9930	1.0081	1.0005	1.0038	1.3678	1.0209	1.2192	1.0102
0.8	0.4	1.1085	1.0118	1.0646	1.0069	0.9893	1.0090	1.0067	1.0001	1.3847	1.0183	1.1988	1.0227

Notes: (i) This table reports the results of the simulations (social costs, costs and payments) under each mechanism relative to the optimal mechanism. It assumes that $\zeta = 1, 800$. MTQs yield the lowest social costs for all parameters considered. Technology separation yields the lowest payments, but higher costs. Technology neutrality yields the lowest costs, but higher payments. Technology separation outperforms technology neutrality except for all cases with $\lambda = 0$ and the case with $\lambda = 0.2$ and $\rho = -0.8$. (ii) Note that, for cases with $\lambda = 0$, payments under neutrality, banding and MTQs are higher than under the optimal mechanism. This is driven by the fact that the costs of the two technologies do not exactly coincide at the margin in the presence of step-wise (average cost) functions. Since the optimal mechanism pays technology-specific prices, it thus saves this price difference.

favour the technology-neutral over the technology-specific approach, when the cost of public funds is sufficiently low and the cost correlation is sufficiently negative.

The table also confirms that banding outperforms technology neutrality, but it only outperforms technology separation when λ is small (i.e., for all three values of ρ , the technology-specific approach gives rise to lower social costs when $\lambda = 0.4$). In any event, all three formats are outperformed by MTQs. The reason is simple: separating technologies for the more extreme cost realisations is effective in containing payments when this is most needed, while allowing for neutrality when cost shocks make technologies more symmetric is effective in avoiding cost inefficiencies. Indeed, social costs depart only slightly from those under optimality. We have performed robustness checks by choosing other parameter sets and this conclusion remains intact (see Appendix D). In other contexts, if technology asymmetries are milder, or if the slope of the solar curve becomes flatter while that of wind becomes steeper, the results could well change in favour of technology banding.

In sum, the results of our simulations suggest that Spain's novel MTQ design might have been a good choice over alternative formats, given the current state of the technologies; or, more precisely, if the (relative) state of the technologies that was expected for 2021 did not differ much from the one observed in 2019 (as the data reported by IRENA (2020) indeed suggests).

7. Conclusions

Our paper analyses an issue that is at the heart of a successful energy transition, namely, how to optimally procure low carbon technologies at least cost for society. In particular, we have shed light on whether and when to favour a technology-neutral versus a technology-specific approach, and whether and when to do so under price or quantity regulation. Regulators worldwide have favoured one approach or another without there being a more formal analysis of the trade-offs involved; particularly so when one takes into account the budget constraint faced by regulators. We have shown that there does not exist a one-size fits all solution: the preferred instrument should be chosen on a case-by-case basis, depending on the characteristics of the technologies and the information available to the regulator.

We have shown that the comparison of a technology-neutral versus a technology-specific approach is faced with a fundamental trade-off. By allowing quantities to adjust to cost shocks, the technology-neutral approach achieves cost efficiency at the cost of leaving high rents with inframarginal producers. In contrast, the technology-specific approach sacrifices cost efficiency in order to reduce those rents. In doing so, it also exploits the benefits that accrue from the (possibly, imperfect) substitutability across technologies. Therefore, whether one approach dominates over the other depends on the specifics of each case.

In particular, technology-specific auctions tend to dominate technology-neutral auctions when technologies are fairly asymmetric—as in our simulation exercise based on detailed information from solar and wind investments in Spain—and the costs of public funds are large, which is when the rent extraction motive is stronger. The opposite is true when cost uncertainty is large and cost shocks are negatively correlated, which is when the concerns for cost efficiency matter most. A low degree of substitutability across technologies further favours the technology-specific approach.

The extremes of technology neutrality and separation can be improved by considering hybrid designs that introduce either technology banding or MTQs. Even in the case of perfect substitutes, technology neutrality is always dominated by technology banding, which in turn dominates

technology separation, but only when cost shocks are sufficiently negatively correlated. Setting MTQs dominates both technology neutrality and separation, and might also dominate banding if the cost correlation is positive and large.

Last, while technology-specific prices always dominate a technology-neutral price, the comparison with the quantity instruments again depends on parameter values. A convex cost curve relative to the benefit curve favours the price approach, while small cost asymmetries across technologies and low costs of public funds tend to further favour the choice of a single quantity target over the choice of technology-specific prices.

We believe that the procurement of green technologies is a most natural application of our analysis. Beyond the reasons we already discussed in the introduction, we want to conclude by highlighting a key fact: namely, in the energy sector, there is typically a single principal (e.g., the national or the supranational regulator). This means that, if she opts for technology separation, she decides on the quantity targets or the prices for each technology, while internalising the overall effect of such choices on total expected social benefits, costs and payments. Otherwise, in the presence of multiple principals, there would be no reason to expect that the separation of technologies would be done optimally. Indeed, as we have shown in our analysis of procurement auctions, the quantity target of the less efficient technology is distorted upwards in order to reduce total payments, at the expense of increasing the rents left with the inefficient suppliers. For this reason, with two principals, each deciding on a separate auction, the optimal solution would likely not be implemented. Beyond the presence of a single versus multiple principals, the fine-tuning that is needed to implement the optimal solution under technology separation might not always be feasible in practice. Indeed, political economy reasons of all sorts (distributional concerns, the pressure of lobby groups, industrial policy considerations, fairness, etc.) might constrain the implementation of the optimal solution under separation. These reasons might explain why in several settings to which our model applies (notably, emissions markets involving various jurisdictions) the separation solution is doomed to fail even if it is theoretically closer to the optimal solution.

Appendix A. Optimal Mechanism

For completeness, here we derive some of the results reported in the main text for the cases of perfect ($\eta = 1$) and imperfect substitutes ($\eta < 1$). Using Lemma 1, we know that the first best (FB) is the solution to

$$P_t^d(q_t, q_{-t}) = \frac{b + (2 - \eta)\beta q_t/2 + \eta\beta q_{-t} - \lambda\gamma q_t}{1 + \lambda} = P_t^s(q_t) = c_t + \theta_t + \gamma q_t$$

for $t = 1, 2$. Summing the two expressions, and re-arranging, allows us to implicitly define the optimal Q^{FB} ,

$$b + \beta Q^{FB} = (c_1 + c_2 + \theta_1 + \theta_2)(1 + \lambda)/2 + \gamma(1 + 2\lambda)Q^{FB}/2. \quad (A1)$$

Solving yields the total optimal quantity

$$Q^{FB} = \bar{Q} - \hat{Q}(\theta_1, \theta_2), \quad (A2)$$

where

$$\bar{Q} = \frac{2b - (1 + \lambda)(c_1 + c_2)}{(1 + 2\lambda)\gamma - 2\beta},$$

$$\hat{Q}(\theta_1, \theta_2) = \frac{1 + \lambda}{(1 + 2\lambda)\gamma - 2\beta}(\theta_1 + \theta_2).$$

Using these expressions, and solving the above system of equations, yields the optimal quantity allocation across technologies,

$$q_1^{FB} = \bar{Q}/2 + \Phi(\eta, \lambda, \Delta c + \Delta\theta)/2 - \hat{Q}(\theta_1, \theta_2)/2,$$

$$q_2^{FB} = \bar{Q}/2 - \Phi(\eta, \lambda, \Delta c + \Delta\theta)/2 - \hat{Q}(\theta_1, \theta_2)/2,$$

where

$$\Phi(\eta, \lambda, \Delta c + \Delta\theta) = \frac{1 + \lambda}{(1 + 2\lambda)\gamma - 2\beta(1 - \eta)}(\Delta c + \Delta\theta).$$

Appendix B. Proofs

Proof of Lemma 2

Statement (i). The welfare maximising solution under technology neutrality, Q^N , solves

$$E\left[\sum_t \frac{\partial B(q_t, q_{-t})}{\partial q_t} \frac{\partial q_t^N(Q)}{\partial Q}\right]$$

$$= E\left[\sum_t C'_t(q_t^N) \frac{\partial q_t^N(Q)}{\partial Q}\right] + \lambda E\left[\frac{\partial p^N(Q)}{\partial Q} Q^N + p^N(Q) \sum_t \frac{\partial q_t^N(Q)}{\partial Q}\right], \quad (\text{B1})$$

where $p^N(Q)$ is the equilibrium price and $C'_t(q_t^N) = c_t + \theta_t + \gamma q_t^N$.

By construction, (a) $\sum_t q_t^N(Q) = Q$, so (b) $\sum_t \partial q_t^N(Q)/\partial Q = 1$. Moreover, cost minimisation implies that (c) $C'_1(q_1^N) = C'_2(q_2^N) = p^N(Q)$, so (d) $\gamma \partial q_1^N/\partial Q = \gamma \partial q_2^N/\partial Q$ and (e) $\partial p^N(Q)/\partial Q = \gamma \partial q_1^N/\partial Q$. But from (b) and (c) we have $\partial q_1^N/\partial Q = \partial q_2^N/\partial Q = 1/2$, so (e) $\partial p^N(Q)/\partial Q = \gamma/2$. Plugging conditions (c) through (e) into (B1) leads to the first-order condition (FOC)

$$b + \beta Q^N = C'_t(q_t^N)(1 + \lambda) + \lambda \gamma Q^N/2$$

for $t = 1, 2$. Summing the two FOCs, taking expectations and dividing by 2, we arrive at

$$b + \beta Q^N = (1 + \lambda)(c_1 + c_2)/2 + \gamma(1 + 2\lambda)Q^N/2,$$

which is analogous to (A1) in expected terms. It follows that $Q^N = E[Q^{FB}]$ and, using (A2), that $Q^N = Q^{FB}(\theta_1, \theta_2) + \hat{Q}(\theta_1, \theta_2)$.

Statement (ii). Its proof simply follows from comparing (3) and (4) for the optimal mechanism and (7) and (8) for technology neutrality. $\square$

Proof of Lemma 3

Statement (i). The welfare maximising solution under technology separation, q_1^S and q_2^S , solves

$$\frac{\partial B(q_t^S, q_{-t}^S)}{\partial q_t} = E[C'_t(q_t^S)] + \lambda E\left[p_t^S(q_t^S) + \frac{\partial p_t^S(q_t^S)}{\partial q_t} q_t^S\right]$$

for $t = 1, 2$, and where $p_t^S(q_t)$ is the equilibrium price in t 's technology-specific auction and $C'_t(q_t^S) = c_t + \theta_t + \gamma q_t^S$.

Using $C'_t(q_t^S) = p_t^S(q_t)$, summing the two FOCs, taking expectations and dividing by 2, we arrive at

$$b + \beta Q^S = (1 + \lambda)(c_1 + c_2)/2 + \gamma(1 + 2\lambda)Q^S/2,$$

where $Q^S = q_1^S + q_2^S$, which is the same as (A1) in expected terms. It follows that $Q^S = E[Q^{FB}]$ and, using (A2), that $Q^S = Q^{FB}(\theta_1, \theta_2) + \hat{Q}(\theta_1, \theta_2)$.

Statement (ii). Its proof simply follows from comparing (3) and (4) for the optimal mechanism and (11) and (12) for technology separation. $\square$

B.1. Proofs of Propositions 1 and 2

Since the total quantity is the same under the two formats (Lemmas 2 and 3), the welfare comparison only depends on the comparison across formats in terms of payments and costs. The comparison then follows immediately using (18), (19), (20), and $\Phi(\eta, \lambda, \Delta c)$ and $\Phi(1, 0, \Delta c)$, as defined in (17). The proof of Proposition 1 is a special case, with $\eta = 1$.

Proof of Proposition 3

From (24) and (25) in the main text, we can write the expected welfare under banding as

$$W_q^B = \bar{W}_q^B(\alpha^B, Q^B) - \frac{\sigma^2[\rho(1 + (\alpha^B)^2) - 2\alpha^B + \lambda(1 + \rho)(\alpha^B - 1)^2]}{\gamma(1 + \alpha^B)^2},$$

where $\bar{W}_q^B(Q^B, \alpha^B)$ corresponds to the deterministic part of the welfare expression and $\{\alpha^B, Q^B\} = \arg \max_{\alpha, Q} W_q^B(\alpha, Q; \sigma, \rho)$.

According to Lemma 2, $\bar{W}_q^B(\alpha, Q)$ is a concave function that reaches its peak when $\alpha = p_2^S/p_1^S = \alpha^B(\sigma = 0) \equiv \alpha_0^B$ and $Q = q_1^S + q_2^S = Q^B(\sigma = 0) \equiv Q_0^B$, that is, when $\bar{W}_q^B(\alpha_0^B, Q_0^B) = W_q^S$ (this is because W_q^S is invariant to shocks). When $\sigma > 0$, $\bar{W}_q^B(\alpha^B, Q^B) < \bar{W}_q^B(\alpha_0^B, Q_0^B)$ and the first-order condition that solves for $\alpha^B(\sigma > 0)$ is given by

$$\frac{\partial \bar{W}_q^B(\alpha^B, Q^B)}{\partial \alpha} - \frac{2\sigma^2(1 + \rho)(2\lambda + 1)(\alpha^B - 1)}{\gamma(1 + \alpha^B)^3} = 0.$$

Since the second term is negative, $\partial \bar{W}_q^B(\alpha^B, Q^B)/\partial \alpha > 0$ and, therefore, $\alpha^B < \alpha_0^B$.

Conditions (i) $\bar{W}_q^B(\alpha^B, Q^B) < \bar{W}_q^B(\alpha_0^B, Q_0^B) = W_q^S$ and (ii) $\alpha^B < \alpha_0^B$ act in different directions as to their impacts on $\bar{\rho}$. While (i) calls for a lower $\bar{\rho}$, (ii) calls for a higher one. To see which effect dominates, take the condition that defines $\bar{\rho}$, i.e.,

$$\bar{W}_q^B(\alpha^B, Q^B) - \frac{\sigma^2[\bar{\rho}(1 + (\alpha^B)^2) - 2\alpha^B + \lambda(1 + \bar{\rho})(\alpha^B - 1)^2]}{\gamma(1 + \alpha^B)^2} = W_q^S, \quad (B2)$$

and totally differentiate it with respect to σ^2 . Using the envelope theorem yields (note that ρ only enters indirectly in $\bar{W}_q^B$, through its effects on α^B and Q^B),

$$\frac{d\bar{\rho}}{d\sigma} = \frac{-[\bar{\rho}(1 + (\alpha^B)^2) - 2\alpha^B + \lambda(1 + \bar{\rho})(\alpha^B - 1)^2]}{\sigma^2[1 + (\alpha^B)^2 + \lambda(\alpha^B - 1)^2]} > 0.$$

Recall that the numerator is positive because of (i).

It remains to show that $\bar{\rho}$ is bounded away from 1, regardless of σ . We proceed by contradiction. If $\bar{\rho}$ were to approach unity for some value of σ then, from (B2), we would obtain $\bar{W}_q^B(\alpha^B, Q^B) > W_q^S$, a contradiction. $\square$

Proof of Proposition 5

Let p_1^* and p_2^* be the optimal posted prices, leading to equilibrium quantities

$$q_t(p_t^*) = (p_t - c_t - \theta_t)/\gamma$$

and welfare

$$W_p^S = E \left[bQ_p + \frac{\beta}{2}(Q_p)^2 - \sum_{t=1,2} \left\{ (c_t + \theta_t)q_t(\cdot) - \frac{\gamma}{2}(q_t(\cdot))^2 - \lambda p_t^* q_t(\cdot) \right\} \right],$$

where $Q_p = q_1(p_1^*) + q_2(p_2^*)$. For the same reasons that the deterministic component under the (optimal) price design in Weitzman (1974) is equal to the deterministic component under the (optimal) quantity design, here the deterministic component of W_p^S is equal to W_q^S ; therefore, ΔW_{pq}^S is simply the stochastic component, which is

$$\frac{\beta}{2\gamma^2} E[(\theta_1 + \theta_2)^2] + \frac{1}{2\gamma} (E[\theta_2^2] + E[\theta_1^2]),$$

or (27) in the text. $\square$

Appendix C. Market Power

In the main text, we have assumed that suppliers behave competitively by offering their units at marginal cost. In this section, we revisit our previous analysis of technology-neutral and technology-specific auctions by adding market power to the model.⁴² We stick to the assumption of perfectly substitutable technologies, $\eta = 1$. Furthermore, since we do not want to introduce asymmetries across technologies, we assume a symmetric market structure for both, with one dominant firm (d) controlling a share ω of each unit, while the remaining share, $1 - \omega$, belongs to a fringe of competitive firms (f). Aggregate costs remain unchanged, while the costs faced by the dominant firm and the fringe now differ. In particular, the costs for each $i = d, f$ are given by

$$C_{it}(q_{it}) = (c_t + \theta_t)q_{it} + \frac{1}{2} \frac{\gamma}{\omega_i} q_{it}^2$$

with $\omega_d = \omega$ and $\omega_f = 1 - \omega$. Accordingly, the higher ω , the more efficient is the dominant firm relative to the fringe and the stronger is its market power.⁴³

⁴² Similar conclusions would be obtained if we also compared these to banding.

⁴³ The presence of a dominant firm opens up the door for non-linear mechanisms; for instance, they could involve menus with quantity discounts (premia, in this case). The extent to which our finding below (Proposition 1)—i.e., that market power favours the neutral approach over the specific one—remains in the context of non-linear menus will depend,

While the fringe behaves competitively, the dominant firm sets prices in order to maximise its profits over the residual demand. Under technology neutrality, the market clearing price now becomes

$$p^N(\theta_1, \theta_2) = \frac{c_1 + c_2 + \theta_1 + \theta_2}{2} + \frac{\gamma}{1 - \omega^2} \frac{Q^N}{2},$$

which corresponds to our previous solution for $\omega = 0$, (6). As ω goes up, the slope of the price equation becomes steeper.

The resulting expected allocation across firms is

$$E[q_d^N] = \frac{\omega}{1 + \omega} Q^N < E[q_f^N] = \frac{1}{1 + \omega} Q^N,$$

with both firms ex post allocating their production across technologies in order to equalise their marginal costs. The market share of the dominant firm is smaller as it withholds output to push prices up.

Likewise, under technology-specific auctions, the market clearing price becomes, for $t = 1, 2$,

$$p_t^S(\theta_t) = c_t + \theta_t + \frac{\gamma}{1 - \omega^2} q_t,$$

and the resulting allocation across firms is

$$q_{dt}^S = \frac{\omega}{1 + \omega} q_t^S < q_{ft}^S = \frac{1}{1 + \omega} q_t^S.$$

Similarly to our first lemma, Lemma C.1 below compares the quantity choices under technology-neutral and technology-specific auctions in the presence of market power.

LEMMA C.1. *For all ω , the optimal total quantities in a technology-neutral auction and technology-specific auctions are the same, i.e., $Q^N(\omega) = Q^S(\omega)$, but the expected quantities allocated to each technology are not: $q_1^S(\omega) < E[q_1^N(\omega)]$ and $q_2^S(\omega) > E[q_2^N(\omega)]$. In turn, $Q^N(\omega)$ and $Q^S(\omega)$ are decreasing in ω , and the allocative distortions $E[q_1^N(\omega)] - q_1^S(\omega)$ and $q_2^S(\omega) - E[q_2^N(\omega)]$ are increasing in ω .*

PROOF. See below. □

As in perfectly competitive auctions, the regulator chooses the same aggregate quantity across the two approaches, but distorts the technology-specific targets from the ex ante efficient solution. Interestingly, market power adds new twists. First, in the presence of market power, increasing the total quantity involves higher marginal costs, given that market power distorts the quantity allocation across firms. It also increases payments more, as market power results in higher prices and makes the price curve steeper. Since the marginal benefits are unchanged, it follows that, when the total quantity procured is lower, the greater the degree of market power.

Second, market power affects the distortion in the technology-specific targets. Two forces are moving in opposite directions. Because the price curves are steeper, marginally moving quantity from the low-cost to the high-cost technology reduces payments relatively more than in the absence of market power. However, because market power distorts the quantity allocation across firms, distorting the allocation across technologies increases costs more than in the absence of market power. The first effect dominates, however, leading to more quantity distortion across technologies as market power goes up.

among others, on whether menus' incentive compatibility constraints are cheaper to handle under separation than under neutrality. However, exploring this possibility in detail is out of the scope of this paper.

The comparison between technology neutrality and separation still reflects a rent-efficiency trade-off, with the former being more effective at reducing costs and the latter being more effective at containing payments. Market power affects these two objectives, increasing costs and payments under both approaches. However, the comparison is tilted in favour of technology neutrality. The reason is two-fold. First, through the effect of market power on the quantity distortion, the cost increase is higher under technology separation than under technology neutrality. And second, separation is increasingly less effective in reducing overall payments as market power goes up. This is stated in our last proposition.

PROPOSITION C.1. *Market power reduces welfare under both approaches, but the welfare reduction is greater under technology-specific auctions, i.e., $W_q^N - W_q^S$ is increasing in ω .*

PROOF. See below. □

To gain some intuition, consider the extreme case of a monopolist facing either one or two inelastic quantity targets. In either case, the monopolist would charge the highest possible price, fully offsetting the possibility to reduce payments through separation. Hence, expected payments would be equal under both types of auction. However, unlike technology separation, technology neutrality would allow the monopolist to freely allocate its production across technologies. As this reduces total costs, the presence of a monopolist does not hurt welfare as much under technology neutrality as under separation. For not so extreme degrees of market power, the technology-specific approach may still dominate technology neutrality, but the range of parameter values for which this is the case is narrower than in Proposition 1.

Proof of Lemma C.1

To show that $Q^N(\omega) = Q^S(\omega)$, we start by considering the FOC that solves for $Q^N(\omega)$,

$$\begin{aligned} \frac{\partial B(Q^N)}{\partial Q} = E \left[\sum_{i=f,d} \sum_{t=1,2} \frac{\partial C_{ti}(q_{ti}^N)}{\partial q_{ti}} \frac{\partial q_{ti}^N(Q)}{\partial Q} \right] \\ - \lambda E \left[\frac{\partial p^N(Q)}{\partial Q} Q^N + p^N(Q) \sum_{i=f,d} \sum_{t=1,2} \frac{\partial q_{ti}^N(Q)}{\partial Q} \right], \end{aligned} \quad (C1)$$

where $p^N(Q)$ is the equilibrium price and $\partial C_{ti}(q_{ti}^N)/\partial q_{ti} = c_t + \theta_t + \gamma q_{ti}^N/\omega_t$.

Expression (C1) can be simplified using several conditions that must hold in equilibrium, such as the balance condition (i) $Q = \sum_i \sum_t q_{ti}^N(Q)$ and the cost-minimising condition (ii) $\partial C_{ti}(q_{ti}^N(Q))/\partial q_{ti} = \partial C_{-ti}(q_{-ti}^N(Q))/\partial q_{-ti}$ for $t = 1, 2$ and $i = f, d$. Totally differentiating these two conditions with respect to Q adds two further conditions: (iii) $1 = \sum_i \sum_t \partial q_{ti}^N(Q)/\partial Q$ and (iv) $\partial q_{li}^N(Q)/\partial Q = \partial q_{2i}^N(Q)/\partial Q$ for $i = f, d$, respectively. In addition, we have the fringe's price-taking condition (v) $p^N(Q) = \partial C_{tf}(q_{tf}^N(Q))/\partial q_{tf}$ for $t = 1, 2$, which, in turn, leads to condition (vi): $\partial p^N(Q)/\partial Q = \gamma/(1 - \omega) \times \partial q_{tf}^N(Q)/\partial Q$ for $t = 1, 2$. Finally, we have the dominant firm's profit-maximisation condition

$$\{q_{1d}^N, q_{2d}^N\} = \arg \max \{p^N(Q)(q_{1d}^N + q_{2d}^N) - C_{1d}(q_{1d}^N) - C_{2d}(q_{2d}^N)\}, \quad (C2)$$

subject to (i) and (v).

Solving (C2) we arrive at the FOC

$$q_{1d}^N(Q) + q_{1d}^N(Q) - 2(1 - \omega) \left(\frac{1}{1 - \omega} q_{1f}^N(Q) - \frac{1}{\omega} q_{1d}^N(Q) \right) = 0 \quad (C3)$$

for $t = 1, 2$. Totally differentiating (C3) with respect to Q and using (iv) we obtain condition (vii) that reads $\partial q_{1d}^N(Q)/\partial Q = \omega \partial q_{1f}^N(Q)/\partial Q$ for $t = 1, 2$. Furthermore, condition (vii) together with (iii) and (iv) leads to condition (viii): $\partial q_{1f}^N(Q)/\partial Q = 1/2(1 + \omega)$ and $\partial q_{1d}^N(Q)/\partial Q = \omega/2(1 + \omega)$ for $t = 1, 2$. And, since $\partial q_{1i}^N(Q)/\partial Q = \partial q_{2i}^N(Q)/\partial Q$ from (iv), integrating yields

$$q_f^N(Q) = \frac{1}{1 + \omega} Q \quad \text{and} \quad q_d^N(Q) = \frac{\omega}{1 + \omega} Q, \quad (C4)$$

where $q_f^N(Q) = q_{1f}^N(Q) + q_{2f}^N(Q)$ and $q_d^N(Q) = q_{1d}^N(Q) + q_{2d}^N(Q)$. Note that, while $q_i^N(Q)$ is deterministic, $q_{1i}^N(Q)$ and $q_{2i}^N(Q)$ are not.

Substituting (viii) into (C1) yields

$$\frac{\partial B(Q^N)}{\partial Q} = E \left[\frac{\partial C_{1f}(q_{1f}^N)}{\partial q_{1f}} \frac{1}{1 + \omega} + \frac{\partial C_{1d}(q_{1d}^N)}{\partial q_{1d}} \frac{\omega}{1 + \omega} \right] + \lambda E \left[\frac{\partial C_{1f}(q_{1f}^N)}{\partial q_{1f}} + \frac{1}{2} \frac{\gamma}{1 - \omega^2} Q^N \right]$$

for $t = 1, 2$. Summing conditions for $t = 1$ and $t = 2$, using (C4), taking expectations and dividing by 2, we conveniently arrive at

$$\frac{\partial B(Q^N)}{\partial Q} = \frac{1}{2}(1 + \lambda)(c_1 + c_2) + \frac{1}{2}A(\omega)(1 + 2\lambda)\gamma Q^N, \quad (C5)$$

where

$$A(\omega) = \frac{1 + 2\lambda(1 + \omega) + \omega(1 - \omega)}{(1 + 2\lambda)(1 - \omega)(1 + \omega)^2} \quad (C6)$$

with $A(0) = 1$ and $A'(\omega) > 0$ (note that $\text{sign}[A'(\omega)] = \text{sign}[4\lambda(1 + \omega) + 3\omega - \omega^2]$).

Consider now the FOCs that solve for $q_1^S(\omega)$ and $q_2^S(\omega)$,

$$\frac{\partial B(q_1^S + q_2^S)}{\partial q_t} = E \left[\sum_{i=f,d} \frac{\partial C_{ti}(q_{ti}^S)}{\partial q_{ti}} \frac{\partial q_{ti}^S(q_t^S)}{\partial q_{ii}} \right] + \lambda E \left[p_t^S(q_t^S) + \frac{\partial p_t^S(q_t^S)}{\partial q_t} q_t^S \right] \quad (C7)$$

for $t = 1, 2$, where $p_t^S(q_t)$ is the equilibrium price in t 's technology-specific auction and $\partial C_{ti}(q_{ti}^S)/\partial q_{ti} = c_t + \theta_t + \gamma q_{ii}^S/\omega_t$.

Proceeding as above, we obtain

$$q_{1f}^S(q_t^S) = \frac{1}{1 + \omega} q_t^S \quad \text{and} \quad q_{1d}^S(q_t^S) = \frac{\omega}{1 + \omega} q_t^S, \quad (C8)$$

where $q_f^S = q_{1f}^S + q_{2f}^S$ and $q_d^S = q_{1d}^S + q_{2d}^S$. Summing the two FOCs given by (C7), one for each technology, using (C8), taking expectations and dividing by 2, we obtain

$$\frac{\partial B(q_1^S + q_2^S)}{\partial q_t} = \frac{1}{2}(1 + \lambda)(c_1 + c_2) + \frac{1}{2}A(\omega)(1 + 2\lambda)\gamma Q^S, \quad (C9)$$

where $Q^S = q_1^S + q_2^S$.

Looking at (C5) and (C9), it is clear that the two expressions are the same, implying that $Q^N(\omega) = Q^S(\omega)$ for all ω . Furthermore, that $Q^N(\omega)$ and $Q^S(\omega)$ are decreasing in ω follows directly from the concavity of $B(\cdot)$ and $A'(\omega) > 0$.

For the rest of the proof note, after some manipulation, that the presence of market power affects (7), (8), (11) and (12) in the main text as follows

$$\begin{aligned} q_1^N(\omega) &= \frac{1}{2}(Q^N(\omega) + \Phi(0, \Delta c + \Delta\theta)), \\ q_2^N(\omega) &= \frac{1}{2}(Q^N(\omega) - \Phi(0, \Delta c + \Delta\theta)), \\ q_1^S(\omega) &= \frac{1}{2}\left(Q^S(\omega) + \frac{\Phi(\lambda, \Delta c)}{A(\omega)}\right), \\ q_2^S(\omega) &= \frac{1}{2}\left(Q^S(\omega) - \frac{\Phi(\lambda, \Delta c)}{A(\omega)}\right). \end{aligned}$$

Since $\partial[\Phi(\lambda, \Delta c)/A(\omega)]/\partial\omega < 0$ (recall that $A'(\omega) > 0$) and $Q^S(\omega) = Q^N(\omega)$, the distortion

$$E[q_1^N] - q_1^S = q_2^S - E[q_1^N] = (\Phi(0, \Delta c + \Delta\theta) - \Phi(\lambda, \Delta c)/A(\omega))/2$$

is also increasing in ω . $\square$

Proof of Proposition C.1

We want to show that welfare falls with ω under both approaches, but more so under the technology-specific approach. Using (15) in the main text and the expressions in Lemma 3, we can compute, after some algebra, the difference in expected costs as

$$\Delta C^{SN}(\omega) \equiv E[C^S(Q^S(\omega))] - E[C^N(Q^N(\omega))] = \Delta C^{SN}(0) + \Psi(\omega) > 0,$$

where

$$\Psi(\omega) = \frac{\omega^3 \gamma [\Phi(\lambda, \Delta c)]^2}{4(1 + \omega)^2(1 - \omega)} > 0$$

with $\Psi(0) = 0$ and $\Psi'(\omega) > 0$. This shows that, as we increase market power, the cost difference also goes up due to the further allocative distortion under separation.

Similarly, and following (14), the difference in payments can be written as

$$\Delta T^{SN}(\omega) \equiv E[T^S(Q^S(\omega))] - E[T^N(Q^N(\omega))] = \Delta T^{SN}(0) \Upsilon(\omega) < 0,$$

where

$$\Upsilon(\omega) = \frac{1}{\lambda A(\omega)} \left[1 + 2\lambda - \frac{1 + \lambda}{(1 - \omega^2)A(\omega)} \right] > 0$$

with $\Upsilon(0) = 1$ and $A(\omega)$ given by (C6). Since $A'(\omega) > 0$ and $\partial[(1 - \omega^2)A(\omega)]/\partial\omega < 0$, $\Upsilon'(\omega) < 0$ in the relevant range, that is, when $\Upsilon(\omega) > 0$. And since $\Delta T^{SN}(0) < 0$, then $\Delta T^{SN}(\omega)$ is increasing in ω , reducing the advantage of separation from a payment perspective. It follows that welfare decreases more with ω under separation than under neutrality. $\square$

Appendix D. Additional Simulation Results

See below.

Table D.1. *Simulation Results Relative to the Optimal Mechanism (Smaller Uncertainty).*

ρ	λ	Social Costs				Costs				Payments			
		Neutral	Specific	Banding	MTQ	Neutral	Specific	Banding	MTQ	Neutral	Specific	Banding	MTQ
-0.8	0	1.0000	1.0420	1.0000	1.0000	1.0000	1.0420	1.0000	1.0000	1.0596	0.7646	1.0596	1.0596
-0.8	0.2	1.0685	1.0310	1.0477	1.0066	0.9932	1.0384	1.0065	1.0058	1.4087	0.9980	1.2336	1.0099
-0.8	0.4	1.1219	1.0271	1.0774	1.0082	0.9887	1.0336	1.0019	1.0012	1.4286	1.0121	1.2510	1.0241
0	0	1.0000	1.0200	1.0000	1.0000	1.0000	1.0200	1.0000	1.0000	1.0337	0.7621	1.0337	1.0337
0	0.2	1.0607	1.0119	1.0365	1.0035	0.9922	1.0121	1.0056	1.0031	1.3717	1.0112	1.1770	1.0049
0	0.4	1.1095	1.0126	1.0592	1.0049	0.9888	1.0118	1.0021	0.9997	1.3879	1.0145	1.1909	1.0168
0.8	0	1.0000	1.0009	1.0000	1.0000	1.0000	1.0009	1.0000	1.0000	1.0037	0.8683	1.0037	1.0037
0.8	0.2	1.0540	1.0024	1.0204	1.0013	0.9909	1.0006	1.0020	0.9999	1.3416	1.0103	1.1040	1.0080
0.8	0.4	1.0989	1.0033	1.0341	1.0026	0.9871	1.0014	0.9981	1.0007	1.3577	1.0076	1.1173	1.0070

Notes: This table reports simulation results when the regulator faces smaller uncertainty ($\zeta = 900$). We report results for social costs, costs and payments under each mechanism relative to the optimal mechanism.

Universidad Carlos III de Madrid, Spain

CEPR, UK

Pontificia Universidad Católica de Chile, Chile

ISCI, Chile

Additional Supporting Information may be found in the online version of this article:

Replication Package

References

- Borenstein, S., Bushnell, J., Wolak, F. and Zaragoza-Watkins, M. (2019). 'Expecting the unexpected: Emissions uncertainty and environmental market design', *American Economic Review*, vol. 109(11), pp. 3953–77.
- Bulow, J. and Roberts, J. (1989). 'The simple economics of optimal auctions', *Journal of Political Economy*, vol. 97(5), pp. 1060–90.
- European Commission. (2014). 'Communication from the commission: Guidelines on state aid for environmental protection and energy 2014–2020', *Official Journal of the European Union*, 200, pp. 1–55.
- Fabra, N., von der Fehr, N.H. and Harbord, D. (2006). 'Designing electricity auctions', *RAND Journal of Economics*, vol. 37(1), pp. 23–46.
- Fowlie, M. (2017). 'The renewable energy auction revolution, energy institute blog', <https://energyathaas.wordpress.com/2017/08/07/the-renewable-energy-auction-revolution/> (last accessed: 1 November 2022).
- Frost, T., Govier, N. and Horn, T. (2015). 'Innovations in the bank's provision of liquidity insurance via indexed long-term repo (ILTR) operations', *Quarterly Bulletin Q2*, Bank of England.
- Google. (2020). 'Realizing a carbon free future: Google's third decade of climate action', <https://www.gstatic.com/gumdrop/sustainability/carbon-free-by-2030.pdf> (last accessed: 1 November 2022).
- IRENA. (2019). 'Renewable energy auctions: Status and trends beyond price', Abu Dhabi: International Renewable Energy Agency, https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2019/Jun/IRENA_Auctions_beyond_price_2019_findings.pdf (last accessed: 1 November 2022).
- IRENA. (2020). 'Renewable power generation costs in 2019', Abu Dhabi: International Renewable Energy Agency, <https://www.irena.org/publications/2020/Jun/Renewable-Power-Costs-in-2019> (last accessed: 1 November 2022).
- Jones, M. and Pakkai, V. (2019). 'Technology-neutral auctions—myth or reality?', ICIS Energy Connections Blog, <https://www.icis.com/energy-connections/2019/03/technology-neutral-auctions-myth-or-reality/> (last accessed: 1 November 2022).
- Klemperer, P. (2010). 'The product-mix auction: A new auction design for differentiated goods', *Journal of the European Economic Association*, vol. 8(2–3), pp. 526–36.
- Laffont, J.-J. (2005). *Regulation and Development*, Cambridge: Cambridge University Press.
- Laffont, J.-J. and Martimort, D. (2002). *The Theory of Incentives. The Principal Agent Model*, Princeton, NJ: Princeton University Press.
- Laffont, J.-J. and Tirole, J. (1993). *A Theory of Incentives in Procurement and Regulation*, Cambridge, MA: MIT Press.
- Laffont, J.-J. and Tirole, J. (1996). 'Pollution permits and compliance strategies', *Journal of Public Economics*, vol. 62(1–2), pp. 85–125.
- Lamp, S., Samano, M. and Tiedemann, S. (2021). 'Learning in multi-unit auctions for renewable energy', Working paper, Universidad Carlos III de Madrid.
- Manzano, C. and Vives, (2021). 'Market power and welfare in asymmetric divisible good auctions', *Theoretical Economics*, vol. 16(3), pp. 1095–137.
- Martimort, D. and Sand-Zantman, W. (2016). 'A mechanism design approach to climate-change agreements', *Journal of the European Economic Association*, vol. 14(3), pp. 669–718.
- Mason, C. and Plantinga, A. (2013). 'The additionality problem with offsets: Optimal contracts for carbon sequestration in forests', *Journal of Environmental Economics and Management*, vol. 66(1), pp. 1–14.
- MITECO. (2022). 'Registry of renewable installations in Spain (RIPRE)', Ministerio para la Transición Ecológica y el Reto Demográfico, <https://energia.serviciosmin.gob.es/Pretor/> (last accessed: 1 November 2022).
- Montero, J.-P. (2001). 'Multipollutant markets', *RAND Journal of Economics*, vol. 32(4), pp. 762–74.
- Myerson, R. (1981). 'Optimal auction design', *Mathematics of Operations Research*, vol. 6(1), pp. 58–73.
- Red Eléctrica de España. (2022). 'Electricity generation and consumption data', Electronic Information System from the Spanish Electricity System Operator (ESIOS), <https://www.esios.ree.es/en/generation-and-consumption> (last accessed: 1 November 2022).
- Schmalensee, R. (2012). 'Evaluating policies to increase electricity generation from renewable energy', *Review of Environmental Economics and Policy*, vol. 6(1), pp. 45–64.
- Shubik, M. and Levitan, R. (1980). *Market Structure and Behavior*, Cambridge, MA: Harvard University Press.

- UK Government. (2013). 'Calculating renewable obligation certificates (ROCs)', Department for Business, Energy & Industrial Strategy.
- van Benthem, A. and Kerr, S. (2013). 'Scale and transfers in international emissions offset programs', *Journal of Public Economics*, vol. 107(1), pp. 31–46.
- Weitzman, M. (1974). 'Prices vs. quantities', *The Review of Economic Studies*, vol. 41(4), pp. 477–91.