An open-access obligation is the best way to enhance competition following a Vodafone-Three merger

Comment on possible remedies under Rule 12 of the CMA's rules of procedure for merger, market, and special reference groups

Peter Cramton

Emeritus Professor of Economics at the University of Maryland and International Research Fellow at the Max Planck Institute for Research on Collective Goods

27 September 2024

I am responding to the Notice of Possible Remedies to address competition concerns following a Vodafone-Three merger. I am a researcher and practitioner studying competition issues in mobile communications for over thirty years. I have served as an expert in spectrum auction design for 20 governments and advised 46 bidders in spectrum auctions. Competition issues were an essential element of this work. I have written 52 articles and books on mobile communications, including ways to promote competition. I have advised Ofcom on spectrum auctions several times since 2006 and continue to do so. The work has led to innovations in spectrum auctions adopted globally. My comments today are mine alone and not those of any organization I may be affiliated with.

My expertise is in unconventional merger remedies that intensify competition in the wholesale market by promoting open access. This allows small service providers to enter the market. Enhanced competition in the wholesale market then spills over to the retail market. Consumers and industry benefit.

I introduced this approach when asked to design a remedy for a merger in the electricity sector between EDF and EnBW. The European Commission required a remedy to restore competition following the merger. A traditional divestiture was not possible because of labor-contract issues. Instead, I proposed a virtual divestiture. The merged entity would be required to conduct quarterly auctions for forward energy options for ten years. The reserve price in the auction was zero, guaranteeing that the obligated quantity of forward energy options would sell.

The remedy successfully brought liquidity and forward price information to the wholesale electricity market, thus encouraging the entry of service providers. The remedy was so successful at finding competitive prices that EDF continued to auction forward energy options beyond the ten-year obligation. The remedy's success led to subsequent adoption in many European mergers in Spain, Denmark, Portugal, Germany, Austria, France, and Hungary in the electricity and natural gas sectors. This widespread adoption is a testament to the remedy's effectiveness. See Lawrence M. Ausubel and Peter Cramton, "Virtual power plant auctions," *Utility Policy*, 18, 201-208, 2010. For communications markets, Ofcom held a workshop on this topic, "Ofcom workshop on capacity auctions and open access," on 20 May 2016.

Since 2001, enormous advances in information technology have enabled more powerful yet readily implemented remedies that promote efficient and transparent trade in wholesale markets. These developments were published earlier this month. See Peter Cramton et al., "An open-access market for global communications," *Telecommunications Policy*, 48, 2024. The paper presents an open-access communications market that allows much richer time and location granularity than was possible in 2001. My research team is nearing completion of an open-source commercial platform that would be readily customized to apply the approach to the Vodafone-Three merger. As is typical of non-divestiture remedies, there would be some regulatory oversight to ensure that the obligations of the remedy were met. However, the rewards for enhanced innovation and competition would make it worthwhile for consumers and industry.

The remedy is easy to describe. The merged entity would be obligated to sell a fraction of its capacity in an open-access market for a specified term. Specifically, I will use ten percent (capacity share) and ten years (obligation term) and refer to the merged entity as Vodafone. Vodafone would conduct a competitive procurement to identify the neutral market operator that would conduct the market. The operator would have a well-defined task: conduct an efficient and transparent open-access market for communications capacity. As described in the *Telecommunications Policy* paper, the products traded would be forward communications defined by time and location. The market rules and algorithms are straightforward, building on essential insights in economics and optimization.

Vodafone would be required to sell at least ten percent of its communications capacity at each time and location in the open-access market for ten years. As in the electricity and natural gas mergers, the open-access wholesale market would enable service providers (mobile virtual network operators) to enter and offer consumers innovative service plans. Entry is made much easier with the improved price information and liquidity that efficient and transparent forward trade brings. Market participants can establish positions to manage risk and their needs.

The open-access remedy leverages core economic principles and lets market outcomes flexibly respond to competitive market fundamentals. This reduces overall risk and maximizes the value of the communications capacity.

I have attached the *Telecommunications Policy* and *Utility Policy* papers. I am happy to provide additional information to Vodafone, Three, and the Competition and Markets Authority. Please do not hesitate to reach out if you have any further questions or need more details to make an informed decision.

Telecommunications Policy 48 (2024) 102820

Contents lists available at ScienceDirect

Telecommunications Policy

journal homepage: www.elsevier.com/locate/telpol

An open-access market for global communications☆

Peter Cramton [a,*], Erik Bohlin [b], Simon Brandkamp [c], Jason Dark [d], Darrell Hoy [d], Albert S. Kyle [e], David Malec [e], Axel Ockenfels [c], Chris Wilkens [d]

[a] *University of Maryland USA, and the Max Planck Institute for Research on Collective Goods, Germany*
[b] *Ivey Business School, Canada*
[c] *University of Cologne and the Max Planck Institute for Research on Collective Goods, Germany*
[d] *Cramton Associates, USA*
[e] *University of Maryland, USA*

ARTICLE INFO

Keywords:
Open access
Market design
Satellite communications
Flow trading
Efficient trade
Wholesale market

ABSTRACT

An open-access market design is presented to manage network congestion and optimize network use and value. Open access eliminates the walled-garden approach; instead, it commoditizes communications network capacity while decentralizing access to a transparent wholesale market. It ensures that scarce capacity is put to its best use by providing a platform for efficient trade. The market operates without friction using flow trading. It allows participants to bid persistent piecewise-linear downward-sloping demand curves for portfolios of products, gradually adjusting positions toward targeted needs. Flow trading allows fine granularity of products in time and location, creating complete markets. Liquidity and computational feasibility are maintained despite trading millions of interrelated forward and real-time products. Participants manage risk and adverse price impact through trade-to-target strategies. The market operator clears the market every hour, finding unique prices and quantities that maximize as-bid social welfare. Prices, aggregate quantities, and the slope of the aggregate net demand are public. The market operator observes positions, enabling it to optimize collateral requirements to minimize default risk. Priority pricing is used to manage real-time imbalances. An application of the model is developed for intersatellite wholesale communications with optical (laser-beamed) mesh networks in space, showing several efficiency gains.

1. Introduction

While many communications markets are open to competition, they remain highly concentrated, with significant barriers to entry. Market forces push mobile network operators to consolidate from four to three in many markets and three to two in others. Fixed broadband options often are limited to one or two providers.

Limited competition should not come as a surprise given the considerable fixed cost of building a network. The communications industry quickly consolidated into a "natural" monopoly in the early years of the telephone. Technology advancements have made the monopoly model obsolete, but high fixed costs still limit competition. This is especially true when dominant incumbents have

☆ Support by Rivada Networks and the German Science Foundation through Germany's Excellence Strategy (EXC 2126/1 390838866) is gratefully acknowledged.
* Corresponding author.

https://doi.org/10.1016/j.telpol.2024.102820
Received 17 March 2024; Received in revised form 6 July 2024; Accepted 7 July 2024
Available online 15 July 2024

incentives to discourage competition, as the US history of “Ma Bell” illustrates. The Bell System worked to preserve and extend its monopoly into any device that touched its network. Bell’s behavior eventually led to antitrust lawsuits in the 1970s and Bell’s breakup in the 1980s. The history is similar in other countries (Cave et al., 2019). The road to competitive communications markets is a long one made longer by the actions of dominant incumbents.

Despite open architectures like the internet and the regulated connections between networks (a customer on one network talking with a customer on another), each network is its own walled garden with the network owner as the “chief gardener” in charge of access.[1] Often, that network owner sells connectivity on the network to businesses or other service providers through bilateral transactions. The product offered is dedicated connectivity of opaque quality with static pricing, whether a fixed fee, a per-gigabyte fee, or some combination. These bilateral deals create walled gardens within the network owner’s walled garden. Idle network capacity results in an illusion of scarcity. Value is lost behind the garden walls. Innovation and competition are stifled. Instead of driving value by enabling people and their creativity, the paradigm strands potential.

This paper proposes an open access wholesale market that eliminates walled gardens, building on recent advances in wholesale electricity markets (Cramton et al., 2024).[2] The wholesale market includes a real-time market and a forward market. The open-access market operates without friction using flow trading (Budish et al., 2023). Participants bid persistent piecewise-linear downward-sloping net demand curves for portfolios of products. A market operator may be designated by the market players as an independent trading platform, even if the market structure is characterized as an oligopoly. Doing this facilitates network utilization and network rents.

We consider a neutral wholesale network operator that conducts the market to maximize as-bid social welfare. The market rules unambiguously define what is meant by maximizing as-bid social welfare. As-bid means that welfare is optimized with respect to the bids. The market rules state the form of bid expression and the precise optimization that translates bids into prices and quantities. The designated market operator clears the market every hour, finding unique welfare-maximizing prices and quantities. Prices, aggregate quantities, and the slope of the aggregate net demand are made public. The market operator observes positions, enabling it to optimize collateral requirements to minimize default risk.

This open-access market model is applied to intersatellite communication via lasers. A laser-equipped low-earth-orbit mesh network can provide global broadband communications independent of terrestrial fiber networks. Benefits of the technology include 1) global connectivity, 2) low latency since the communications travel 50 percent faster through the atmosphere (299,792 km/s) than fiber (200,000 km/s), and 3) optimized, internet-independent routing for improved latency, reliability, and security. This open-access market complements the rise in broadband satellite competition and terrestrial and satellite communications convergence. However, satellite networks come with one limitation. Capacity is much more constrained than terrestrial fiber networks. The limiting factor is earth-satellite throughput; satellite-to-satellite laser throughput is not expected to be an additional bottleneck (laser throughput is between 20 and 100 Gbps). Maximum earth-satellite throughput varies by satellite and atmospheric conditions. A state-of-the-art fiber cable has a throughput of 100,000 Gbps, several orders of magnitude more than earth-satellite throughput. While capacity constraints arise in terrestrial communications networks, network congestion is much more severe for satellite communications, so its management is critical. Revenue management is also essential given the large network-deployment fixed costs. With open access, revenue management is addressed with the supplier offering supply at prices that exceed marginal cost, much as an airline offers seats, especially in premium classes, at prices above marginal cost. This paper addresses the emerging intersatellite communication market by developing an open access model.

Developing an open-access market would be irrelevant if no supplier had an incentive to adopt it. We argue that adopting and committing to the open-access market serves a supplier’s (satellite provider’s) interest. Doing so maximizes the network’s value so long as the communications market is competitive on the buy side. There are many communications buyers; even the largest buyer purchases a small fraction of the total capacity. Each supplier has an interest in maximizing the value its supply brings to buyers. The open access market creates value by managing congestion efficiently, instead of rationing quantity.

Questions can be raised about whether a satellite provider can be expected to commit to an open-access market. The key is the delegation of network operation to an operator whose mission is open access and who is constrained to operate under the open access rules. This is practically accomplished with governance rules that make it difficult to alter the core tenets of open access. Commitments of this sort are commonplace and especially easy to enforce when the tenets are broadly consistent with the supplier’s interests, such as trading platforms in financial markets. We view this as a reasonable assumption, and with this, the oligopoly structure of the satellite supply side is irrelevant.

Our methodology is market design. We start with the market’s objective and recognize the potential market failures that must be mitigated to achieve the objective. Unlike mechanism design (Myerson, 2008) and industrial organization theory (Tirole, 2015), which are mathematical tools to characterize theoretical possibilities within an assumed structure, market design is pragmatic and focused on implementation (Chen et al., 2021; Cramton, 2009; Milgrom, 2004; Roth, 2002). We define market rules and a means of preference expression, typically a bidding language, that provides participants with good incentives. The market design approach maximizes “as-bid” social welfare, recognizing that truthful bidding is only an approximation. We sacrifice rigorous welfare theorems and instead are

[1] The height of the garden walls varies among countries, depending primarily on regulatory decisions and the legal framework. For example, European mobile communications are closer to open access than North America. Europe’s “roam like your home” policy is the best example.

[2] The day-ahead and real-time markets of US wholesale electricity markets successfully illustrate the open-access approach described here. Both markets have rich time and location granularity with tens of thousands of interrelated products that simultaneously clear. Market operators are responsible for maximizing social value, and some markets utilize a version of flow trading, Cramton (2017); Cramton and Ockenfels (2024).

content with a market that works well in practice.[3]

The paper presents a market design for an open access global communications market. It starts by discussing satellite networks, the market structure, and the open access innovation. Objectives and the benefits of the open-access approach are discussed. The market participants, governance, the role of forward trading, and the distinction between real-time and forward trading are examined. An outline of the market rules, the granularity of products, and the approach to efficient trade are discussed. The simplicity of participation in a market with fine product granularity is explained. Before concluding, the paper discusses liquidity, counterparty risk, flexibility, and competition.

2. Satellite global networks

Space-based optical mesh networks offer high-performing global communications. At least three networks are proposed with intersatellite laser technology: 1) Amazon, 2) Starlink, and 3) Rivada. Starlink performed successful tests in September 2023; Amazon had successful tests in November 2023; Rivada plans its first launch in 2024. There are essential differences between these constellations: Starlink and Amazon are internet service providers, connecting an end-user's dish at one end to an internet gateway. They are using laser links to extend the distance from the gateway from which they can offer service. In contrast, Rivada's planned optical mesh network is currently the only one designed to route traffic across its constellation from end to end. It is not intended to provide consumer-grade internet access. Rather, it is to function as an internet backbone in space—an Outernet.

These differences are important. If the Rivada model proves successful, then we would expect others to move in this direction. Laser-linked low-earth-orbit constellations will evolve in the direction of true orbital mesh networks over time. Low-earth-orbit architectures have three classifications. A Type I constellation like Starlink's current offering provides a last-mile gateway. The satellite is used only for earth-satellite-earth communications to expand connectivity to remote locations. All traffic is routed to the nearest internet gateway. A Type II constellation like Amazon's offering and Starlink's next-generation offering provides last-mile connectivity with a laser extension. Earth-satellite-satellite-earth communication is possible, expanding the reach of satellite connections. A type III constellation like Rivada is developing is a true meshed multiprotocol-label-switching network in space. This approach enables end-to-end connectivity over the satellite network with optimized routing and full global reach.

Amazon and Starlink focus primarily on the retail consumer market. Rivada concentrates on government and enterprise communications. The retail focus requires much greater capacity, resulting in many more individual users and endpoints on the network. This is one reason that Amazon and Starlink have many more satellites than the smaller Rivada network, which has a smaller number of larger, more capable satellites designed to serve a smaller population of user terminals. The three networks also differ in satellite altitude. Amazon is 610 km above sea level, Starlink is 570 km, and Rivada is 1052 km. A higher altitude implies marginally longer earth-satellite latency, greater visibility, and potentially fewer satellite-satellite hops. See (Pachler et al., 2021) for a detailed capacity estimate based on January 2021 FCC filings. Fiber dominates optical mesh networks in terms of capacity. For densely populated areas, fiber is the preferred choice, at least if the communications are traveling a distance of a few thousand kilometers or less.

Despite these capacity constraints, optical mesh networks in orbit will be critical in meeting burgeoning connectivity demand. First, they provide critical redundancy. When there are power outages, fiber and cellular networks become unreliable. The ground networks often have inadequate backup power. For critical communications, an optical mesh network provides a fallback. All that is needed is battery backup for the receiver and router. Critical communications can continue uninterrupted. Modern routers, such as Eero, automatically switch to backup internet.

Second, optical mesh networks in orbit have a latency advantage for communications over several thousand kilometers or more. Latency-critical applications such as high-frequency trading benefit from low latency. High-frequency traders are in an arms race for speed. The fastest trader gets to pick off stale quotes when market fundamentals change; speed is essential (Budish et al., 2015). Table 1 Theorem 2 gives the distance between the top ten global financial centers. Many are separated by several thousand kilometers.

Tables 2 and 3 show the millisecond time between any pair of top-ten financial centers via optical mesh network or terrestrial fiber. The optical mesh network assumes an additional distance to and from the low-earth-orbit satellite, which takes 7 ms for Rivada (the Amazon and Starlink roundtrip takes about 4 ms). The communications travel near the speed of light in a vacuum (299,792 km/s). The communications in the terrestrial fiber network travel at the speed of light in glass (about 200,000 km/s), two-thirds of the speed in air, due to glass's much lower refractive index.

Table 4 takes the difference between Tables 3 and 2 to yield the time savings of the optical mesh network in orbit. This calculation assumes a straight-line approximation between cities and ignores the number of hops required. The terrestrial fiber network will involve a less straight route and more hops, so these are conservative assumptions. For example, communication between New York and Tokyo can occur with about four hops and a direct polar route in an optical mesh network, such as NY to NY-satellite to North Pole-satellite to Tokyo-satellite to Tokyo. Standard communications may involve more hops and longer latency.

Table 5 summarizes the magnitude of the time savings between each city pair. Over three-quarters of the city pairs have a significant improvement in latency. One millisecond is an eternity for a high-frequency trader. Other communications applications, such as online gaming and video conferencing, also benefit from low latency. Low latency is an essential quality attribute for a segment of demand.

A third and transformative benefit of an optical mesh network is coverage. The supplier's low-earth-orbit constellation can connect

[3] By contrast, mechanism design would make theoretical assumptions and then identify incentive-compatible mechanisms that guarantee truthful bidding is consistent with equilibrium behavior in the assumed framework.

Table 1
Distance between the top-ten global financial centers in kilometers.

City	London	Hong Kong	Singapore	Tokyo	Shanghai	Frankfurt	Paris	Zurich	Sydney
New York	5,570	12,955	15,333	10,849	11,858	6,203	5,837	6,324	15,989
London		9,623	10,848	9,559	9,197	638	344	776	16,994
Hong Kong			2,586	2,880	1,227	9,158	9,626	9,301	7,376
Singapore				5,315	3,806	10,258	10,729	10,294	6,306
Tokyo					1,758	9,332	9,712	9,578	7,827
Shanghai						8,820	9,263	9,010	7,881
Frankfurt							478	304	16,483
Paris								488	16,961
Zurich									16,568

Table 2
Communication time (ms) between the top-ten financial centers via optical mesh network.

City	London	Hong Kong	Singapore	Tokyo	Shanghai	Frankfurt	Paris	Zurich	Sydney
New York	26	50	58	43	47	28	26	28	60
London		39	43	39	38	9	8	10	64
Hong Kong			16	17	11	38	39	38	32
Singapore				25	20	41	43	41	28
Tokyo					13	38	39	39	33
Shanghai						36	38	37	33
Frankfurt							9	8	62
Paris								9	64
Zurich									62

Speed of light in a vacuum = 299,792 km/sec
Round-trip earth to satellite = 7 ms

Table 3
Communication time (ms) between the top-ten financial centers via terrestrial fiber network.

City	London	Hong Kong	Singapore	Tokyo	Shanghai	Frankfurt	Paris	Zurich	Sydney
New York	28	65	77	54	59	31	29	32	80
London		48	54	48	46	3	2	4	85
Hong Kong			13	14	6	46	48	47	37
Singapore				27	19	51	54	51	32
Tokyo					9	47	49	48	39
Shanghai						44	46	45	39
Frankfurt							2	2	82
Paris								2	85
Zurich									83

Speed of light in glass = 200,000 km/sec

two global points at gigabit speeds, with low latency and unmatched security. Typically, each satellite is equipped with four laser intersatellite terminals. The four lasers weave together a mesh network that encircles the globe, enabling direct connections between any point in the space-based network with terrestrial terminals. This separation will improve cyber security and data sovereignty by avoiding terrestrial internet gateways, a critical requirement for many government and enterprise uses. Multiple satellite networks at

Table 4
Communication time savings in milliseconds of optical mesh network.

City	London	Hong Kong	Singapore	Tokyo	Shanghai	Frankfurt	Paris	Zurich	Sydney
New York	2	15	19	11	13	3	3	4	20
London		9	11	9	8	-6	-6	-6	21
Hong Kong			-3	-2	-5	8	9	8	5
Singapore				2	-1	10	11	10	3
Tokyo					-4	9	9	9	6
Shanghai						8	8	8	6
Frankfurt							-6	-7	20
Paris								-6	21
Zurich									21

Distances and times are based on a straight-line approximation. Hop latency and the fiber equivalent hand-offs are ignored.

Table 5
Optical mesh network time savings (ms) between each of the top-ten global financial centers.

	<0 ms	0-2 ms	2-4 ms	4-8 ms	8-16 ms	>16 ms
City pairs	11	1	5	5	17	6
Percent	24%	2%	11%	11%	38%	13%

different altitudes provide valuable redundancy, improving resilience to terrestrial and space threats.

Three proposed optical mesh networks are shown in Table 6. Each differs in altitude—distance from earth—and the number of satellites. Starlink is the most ambitious, with a plan for 42,000 satellites; the table uses 10,000 as a more realistic number for the near term. A lower altitude is preferred for denser networks, improving communication latency over a shorter distance. For example, the straight-line no-hop latency between New York and Los Angeles is 17 ms with Amazon's and Starlink's constellations and 20 ms for Rivada's, slightly faster than terrestrial fiber. Terrestrial fiber dominates for shorter distances, such as New York to Chicago, although terrestrial microwave links are the fastest. For longer distances, such as New York to Tokyo, Rivada's network dominates because the higher altitude reduces the required hops.

The critical limitation of the optical mesh networks is capacity. Optimized routing and pricing are used to manage congestion and maximize capacity. To date, Starlink has relied on traditional static pricing. Rivada plans to incorporate dynamic congestion pricing via an open-access market. Other networks may shift to dynamic pricing, as the efficiency gains are too significant to ignore. However, both Amazon and Starlink are focused on retail broadband, a sector where static pricing and rationing are acceptable. The primary retail use is video streaming, where rationing works reasonably well. Consumers can accept a degradation of video quality from UHD to HD when data rates are limited, and buffers can smooth short-duration congestion. By contrast, when committed to the more demanding requirements of business-to-business and government applications, like Rivada, efficient routing and pricing become essential.

3. Market structure

The economics of optical mesh networks is different from terrestrial fiber. Fiber is a natural monopoly. The dominant expense is laying cable. Once a line is installed, a competing cable is redundant. Despite its capacity virtues, fiber is very expensive to lay. A fiber network covering 80 percent of the US population would cost hundreds of billions of dollars. Only a minority of the global population will be served by fiber.

In contrast, an optical mesh network covering every point on earth can be built for less than 10 billion dollars. The main cost is building and orbiting the satellites. Thus, economies of scale are modest, and building an optical mesh network with twice the capacity costs twice as much. The absence of economies of scale beyond about $5 billion means that the initial entrant, Starlink, cannot deter the entry of others.[4]

Two other fundamental economic trends support the expansion of optical mesh networking in orbit. Launch costs are declining quickly, thanks in no small part to SpaceX, but supported further by other new commercial entrants into a satellite-launch market that

[4] There may be entry barriers beyond scale economies, such as consumer lock-in with contracts or equipment.

Table 6
Coverage area of a satellite constellation by altitude.

	Proposed Project		
	Amazon	Starlink	Rivada
Satellites	3,276	10,000	576
Altitude, km	610	570	1,052
Radius, km	2.9	2.8	3.8
Area, km^2	26	24	46
Coverage	5%	5%	9%

Starlink proposes 42,000 satellites eventually. Coverage is a fraction of the earth's surface area for each satellite.
Coverage is a geometry calculation that increases with altitude because of a longer line of sight.

state actors had formerly dominated. Satellite mass manufacturing is finally coming into its own, promising a future in which modular, mechanized, and cost-effective assembly-line techniques and 3-D printing replace bespoke and costly manufacturing. These market changes are driving down the still-high up-front costs of establishing space-based networks, and so, in turn, improving the economics of adding capacity in orbit.

Expanding penetration also supports entry. Demand for communications has increased as technological advances create new data-hungry applications. Moreover, as these technologies become more powerful, the adoption rates increase. Telephone penetration never exceeded 20 percent. By contrast, mobile phone penetration is 100 percent, and broadband internet penetration is 60 percent and rising. See Fig. 1.

Demand for communications varies by time and place. People are dispersed unequally around the globe, as shown in Fig. 2. More granular depictions of population density would show more heterogeneity. Most people live in densely populated urban areas, concentrating communications demand. By contrast, the supply of satellite communications is distributed much more uniformly globally. Thus, to balance supply and demand, the price of satellite communications must vary by location. Satellite networks need dynamic pricing to balance supply and demand.

Communication demands also vary by time of day and day of week. With static pricing, time-of-use variability is managed with quality reductions. Users may experience poor availability and throughput at peak times. With dynamic pricing, price balances supply and demand, increasing availability and throughput for those users who demand good performance. The customer gets to decide the desired quality of service and is motivated to shift non-urgent communications to periods of lower demand.

Variations of demand over time and place and a constant supply dictate that prices must vary by time and place to balance supply and demand. The alternative is to ration demand during peak periods. Rationing is tolerable only if capacity is rarely constrained. However, the limited capacity of satellite communications implies a need for dynamic pricing (Bobbio et al., 2023). This need is the motivation for the open-access market.

4. The open-access market innovation

The open-access market rests on the principle of efficient pricing. A designated market operator conducts an open-access marketplace that brings fairness, transparency, and ease of access to customers who want to acquire network capacity. A network supplier offers wholesale capacity to customers, primarily governments, multinational enterprises, and communications service providers, including mobile network operators and mobile virtual network operators. Customers can purchase or sell capacity in real-time or on a forward basis, as overseen by the market operator, offering flexibility and the ability to manage risk.

The open-access market is a fundamental shift from connectivity (Mbps) by contract to capacity (GB) on demand when and where needed. Customers will actively influence the wholesale market's design and mechanisms through a dispersed governance structure motivated to optimize the market's performance through continuous improvement. The open-access, dynamically-priced market will spur competition in communications services globally and decentralize the innovative power of those services. This highly efficient market will maximize the network's and its users' potential, with benefits to social and economic development and global security.

Below find more details on the market's objectives, benefits, and operations. For concreteness, this paper presents a blockchain implementation of the primary market, as developed by Rivada, although a traditional implementation works similarly.

5. Objectives of the open-access market

The open-access market derives from basic economic principles. Anyone can participate in the market on open, nondiscriminatory terms. Scarce resources are efficiently priced to balance supply and demand.[5] Open access has four principal objectives.

Efficiency. Open access drives more efficient use of network resources, maximizing value through a dynamic pricing mechanism of a

[5] As mentioned in the introduction, our approach is market design and thus our language is informal in the following sense. The pricing and welfare results are with respect to as-bid preferences not the participants' true preferences. Truthful bidding is only an equilibrium under the assumption of perfect competition.

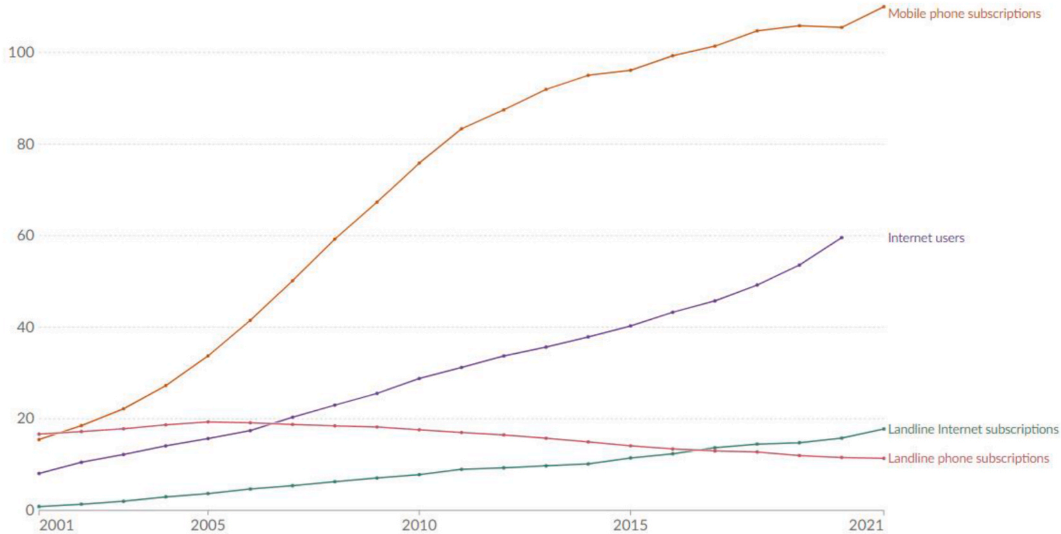

Fig. 1. World adoption of communications technologies, ITU.

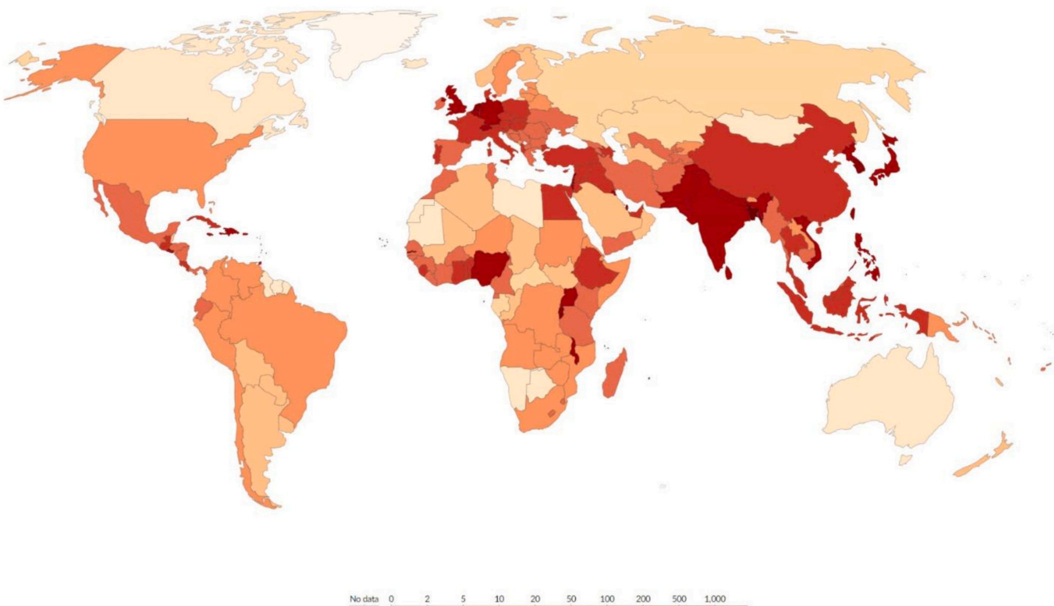

Fig. 2. Population density by country, people per sq km, Our World in Data, 2022.

commoditized unit of network capacity. Market participants can buy or sell capacity on demand in real-time or through a forward market pursuant to their forecasted needs.

Fairness. Open access eliminates discriminatory barriers, providing equal opportunity for everyone to enter the market and buy and sell capacity. The market is indifferent to how the network is used, whether in enterprise applications, competitive wireless or wireline communications, or whether it is resold.

Transparency. Units of network capacity are represented in capacity tokens and traded through blockchain technology. This allows market participants to understand how they are affected by market rules. Prices and aggregate quantities are public as are the rules that determine the mapping of preferences into unique prices and quantities. The market operator monitors participants' positions—their current product holdings. Forward position transparency helps market operators establish optimized collateral requirements and assess market power.

Simplicity. With transparency comes simplicity. By seeing how the market works and having an opportunity to enhance its effectiveness through participation in governance, market participants will clearly understand the market mechanism. Participation is simplified with powerful bid expression and market tools that translate preferences into an effective strategy.

Together, these objectives promote affordability, innovation, and competition in global communications.

6. Benefits of the open-access approach

- The open-access market provides nondiscriminatory access to fungible units of network capacity, allowing supply to respond to demand requirements that vary with location and time.
- Homogenous units of network capacity allow market participants to repurpose efficiently those products to specified heterogeneous needs, allowing for technology-neutral innovation and more competition in communications markets globally.
- The pricing mechanism is transparent and efficient, precisely pricing scarce capacity to balance supply with demand at each time and location. This pricing maximizes the use and value of the network and helps communications providers plan their investments and innovations. It also helps them adjust their plans as their needs evolve with changes in consumer demand.
- The real-time and forward markets allow for granularity in time and location. Granularity encourages a responsive supply of network capacity when and where it is needed most.
- The flow trading approach enables market participants to express preferences and trade in a way consistent with their interests. The method lets market participants satisfy their demands efficiently, create value, and avoid adverse price movements.
- The playing field is level and transparent, with buyers and sellers of capacity having complete visibility into the record of trades.
- Transparency of positions enables all market participants to understand and manage the market's effectiveness and to influence improvements through a novel democratized governance structure.
- Position transparency also allows the market operator to optimize collateral in the forward market to reduce counterparty risk and reduce participants' costs of satisfying collateral requirements.

7. Market participants

Network supplier. The network supplier builds, maintains, and operates the network. The network supplier provides the wholesale service capacity, including its quality, security, and resilience. The network supplier facilitates the open-access wholesale market and supports developing and enhancing the open-access market.

Sellers. As the builder and operator of the wholesale service, the network supplier is the principal seller of capacity on the network. Wholesale customers and other market participants who purchase capacity can, in turn, sell capacity back to the market and thereby become suppliers. By bringing together willing sellers and buyers through a clearing house governed by real-time pricing, the network capacity can be utilized more completely by those who value capacity the most at each time and place.

Buyers. Those seeking access to the network capacity include governments, enterprises, and communications providers, including facilities-based wireless and wireline providers, mobile virtual network operators, and other resellers. Technology companies can leverage global connectivity to create innovative products and solutions.

Independent market operator. The market operator is an independent administrator with a simple mission: "We serve our customers by ensuring secure and reliable communications, efficiently priced in an open-access market." In economic terms, the market operator addresses potential market failures, including incomplete markets, incomplete information, market power, entry barriers, and systemic risk. It also conducts transparent and efficient markets by pricing communications services to maximize as-bid social welfare subject to network constraints. As discussed below, buyers and sellers are eligible to become involved in the governance of the independent market operator (and the open-access market itself) through the rights represented by the blockchain.

Decentralized autonomous organization. The market operator board's oversight of the market is supported by the decentralized autonomous organization, which administers blockchain technology. This governance allows consumers, suppliers, and developers to collaborate and make efficient decisions about the direction and future of the market, which eliminates some of the inefficiency that could plague the independent system operator, as exemplified in electricity markets. The decentralized autonomous organization is run by rules encoded through smart contracts, with all transactions and decision-making processes recorded on the blockchain, thereby enhancing transparency. The decentralized autonomous organization is automated, which improves efficiency and saves costs.

Independent market monitor. An independent panel of experts, which reports to the market operator board, is engaged in objectively evaluating and suggesting continuous improvements to the market operator and its board to enhance the performance of the open-access market.

8. Governance

A critical complement to the market's open-access nature is its democratic governance. Market participants (buyers and sellers) can acquire, hold, and trade governance tokens representing voting and economic rights. These rights include an ownership stake in the market, a right to receive a portion of the market's revenues, the right to elect representatives to the market operator's governing body, and the right to vote on fundamental matters, such as transaction fees, revenue sharing, and treasury allocation. In this way, market participants can influence the design and operation of the market. The governing body of the market operator is the market operator board (Cramton & Doyle, 2017).

9. The role of forward trading

Market participants often wish to procure communications needs in advance to manage risk. The forward market enables them to do so. By buying ahead, the participant can lock in favorable terms. Future purchases also let the participant buy gradually, which reduces trading costs, avoids adverse price movements (Black, 1971; Kyle, 1985; Vayanos, 1999), and ensures that final purchases correspond to needs; uncertainty about these needs is resolved over time.

Forward trade provides information about market fundamentals. This granular price information provides the information and risk-management tools essential to innovation. For the network supplier, the price information is critical for optimizing enhancements to the network, such as capacity additions. Finally, prices encourage resiliency by motivating the participants and market operator to take actions consistent with welfare maximization (Cramton et al., 2024). Forward trade creates a virtuous improvement cycle, as shown in Fig. 3.

10. Distinguishing the physical (real-time) and financial (forward) markets

The market includes a physical real-time market and a financial forward market. Participants use a physical product in real-time by engaging in measured communications through the network. The communications are priced in real-time to balance supply and demand subject to network constraints. The 1-h real-time window means that imbalances typically are insignificant. If real-time rationing is necessary, rationing is limited to throttling or delayed delivery of the regular service type as shown in Fig. 4.

Forward products derived from real-time capacity products are financial; deviations between forward and real-time positions are settled financially. Efficient settlement rests on robust pricing and the elimination of counterparty risk. The market operator manages counterparty risk with collateral obligations that depend on deviations between the participant's forward positions and anticipated needs. The use of smart contracts settled on the blockchain further reduces settlement risk.

11. How will the market work?

All aspects of participation in the open-access market are voluntary, including buying and selling capacity through capacity tokens and acquiring and exercising governance rights through governance tokens. Participants can purchase capacity in the financial forward market and the real-time physical market as price-takers without the need to schedule consumption in advance.

11.1. Products

The primary products traded are the capacity tokens representing capacity on the network, measured in gigabytes, and the governance tokens. These are best thought of as systemwide capacity and governance tokens. The systemwide capacity token is then broken down into communications in gigabytes in an hour, region, and communication type through the open-access auction platform. There are three types—premium, regular, and fast. The communication types differ in their real-time routing optimization. Premium is optimized for reliability and speed; it is nearly never rationed. Regular is optimized for reliability and speed but is rationed as necessary based on network conditions. The premium/regular distinction allows critical communications to be prioritized in the event of excess real-time demand (Chao et al., 1986). Fast is optimized solely for speed, not reliability. Fast is a specialty product tailored to the exacting needs of high-frequency traders and others sending messages where every millisecond matters.

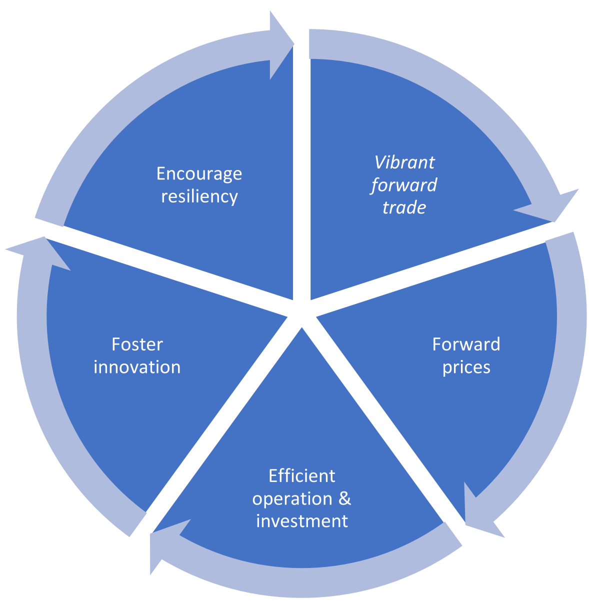

Fig. 3. The virtuous cycle of improvement stemming from forward trade.

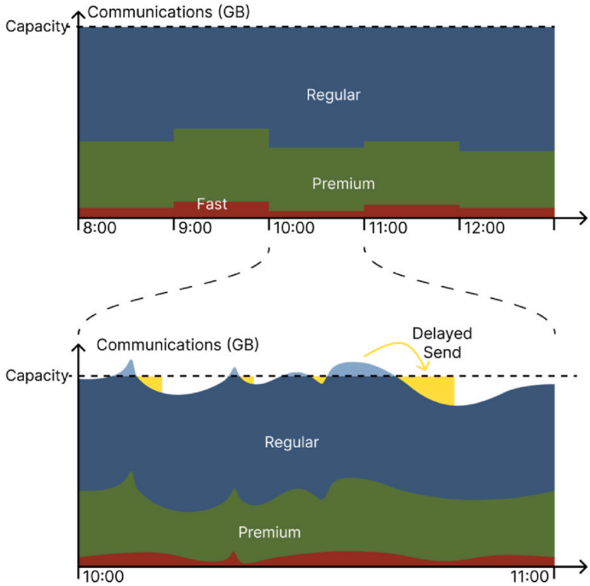

Fig. 4. Regular service may be throttled or delayed when demand exceeds supply in real-time.

The regions are defined as a partition of the globe. A sensible partition does not need to depend on the satellite configuration. Each region is a collection of neighboring areas with similar demand profiles. There are four types of regions:

1. Major metropolitan areas with populations exceeding five million.
2. Other metropolitan areas.
3. Semi-rural areas.
4. Sparsely populated areas.

There are more geographic aggregations as population density and economic activity decline. To simplify, the globe is partitioned into a manageable number of regions. There is little point in differentiating between, say, two locations in the Pacific 100 km from any land mass. In contrast, there are meaningful demand differences between, say, New York City and Atlanta. New York City and Atlanta should be in separate partitions whenever there is economic justification for differential pricing between the two cities. A partition with about five hundred regions seems a good starting point. As the market matures, finer geographic granularity may be desirable.

Regardless of the partition, it is helpful to visualize a hierarchy of partitions of increasing granularity. At the top of the hierarchy is the entire globe.

With 500 regions, there would be $24 \times 500 \times 3 = 36{,}000$ real-time products each day (hours per day $\times$ # of regions $\times$ # of product types). The hourly time granularity for global communications is preferred. Less granular options, such as peak and off-peak, are problematic since these designations are location-dependent. The hourly product has little extra user or system cost. Indeed, finer time granularity is anticipated as the market matures.

11.2. Trading methodology

Flow trading (Budish et al., 2023) allows market participants to adjust user-defined portfolio positions efficiently as information changes over time. Despite fine product granularity, liquidity is not compromised since demand is cleared simultaneously by product independent of portfolio, and trade occurs gradually. Auctions occur every hour for all products. The bidding window starts when prices from the preceding hour are posted, a few minutes after the hour, and lasts until the hour's end. During the bidding window, market participants may adjust their orders. The orders at the end of the bidding window (on the hour) are final. Each order is a piecewise-linear decreasing demand curve, represented by two or more quantity-price pairs. The order also specifies the linear combination of products for which the demand curve applies. The price is in $/GB. Thus, quantity represents the rate of trade—the quantity in gigabytes that trades over the unit of time at a particular location. Market participants either upload their orders or enter them directly into the auction platform. Changes to orders are allowed until the bidding window closes on the hour. Orders that are not valid are rejected. On the hour, the auction platform processes the final orders and determines the prices and quantities that maximize as-bid social welfare.

Note that in a blockchain implementation, the currency is capacity tokens. The token can be exchanged for any other crypto or fiat currency on exchanges outside the open-access market.

Mathematically, the form of preference expression guarantees unique quantities (Budish et al., 2023, Theorem 1). Prices exist but may not be unique (Budish et al., 2023, Theorem 2). However, unique clearing prices result with an intuitive tie-breaking rule: If a product has multiple clearing prices, the price closest to the preceding clearing price is selected. The auction platform revises each participant's position based on the quantities implied by the prices. Each participant can view and download prices and the revised position anytime during the bidding window. This process repeats every hour. Orders persist until changed or canceled. Thus, if the user wants to maintain the same preferences because nothing has changed, then the user does nothing. The same orders will continue to be processed every hour until the user submits a change. The 1-h clearing frequency is a parameter of the market. Faster frequencies, every minute or second, are possible.

11.3. Settlement and collateral

Every hour, the auction platform updates the settlement for each participant. If the user's excess collateral falls below a warning trigger, the user is warned. If the user's excess collateral falls below zero, future trades that would increase the participant's collateral requirement are not allowed. For every order, the portion that would shift the user to a less balanced position are cancelled.

Collateral requirements are based on a to-be-developed optimization. This approach will maximize market stability while minimizing unnecessary capital commitments from participants. The key inputs in determining collateral are the participant's current position and the participant's expected position. The participant estimates the latter and reports it to the market operator. Excessive deviations between reported estimates and realized positions increase the participant's collateral requirements.

Once per week, or more frequently if desired, the accumulated settlement over the preceding 7×24 h is reported to each market participant. Consistent with the weekly accumulation, an automatic transfer to or from the participant's capacity token account is made. If the transfer fails, then the participant has 24 h to resolve the issue. After 24 h, any payment due is taken from the participant's collateral account, and the user is prohibited from further trades that would put the participant in a less balanced position.

11.4. Transparency

The auction platform publicly posts prices when the computation is complete, usually within a few minutes of the end of the hour. Each participant also learns its revised position. The platform lets participants view and download prices, a participant's current position, and the most recent trade rates.

11.5. Market operator

The market operator oversees the real-time and forward market functions through the blockchain to ensure smooth operation. The market operator provides monthly, quarterly, and annual reports on the market performance to the market operator board. Participants have direct access to the blockchain technology from which those reports are generated. The market monitor also studies and discusses the market's performance in state-of-the-market reports.

12. Product granularity

Flow trading allows finer product granularity. The reason is that participants place persistent portfolio orders that induce a smooth trade flow among all products. First, look at the forward prices to understand how this would work. Figs. 5–7 show the yearly, monthly, and hourly hypothetical forward prices for the New York premium capacity during a weekday.[6]

There are many prices, but they are readily understood by the eye and analyzed with computer modeling. The example above illustrates time-of-day and day-of-week price impacts and greater volatility of prices, the closer to real-time. Prices are updated hourly when the market clears.

Participants trade forward products up to five years ahead and adjust positions by hour.

A simple and effective flow trading strategy is trade-to-target, illustrated in Fig. 8. Participants state their target and the rate at which they want to move toward the target. For example, a service provider might set its target to its capacity needs, increasing linearly from zero to expected demand, moving from 5 years ahead to real-time. With flow trading, the participant specifies the rate at which it desires to make this adjustment as a function of price. A participant wants to buy more quickly when the price is low and sell more quickly when the price is high. This is expressed as a linear net demand curve for each product. The participant's urgency to trade also depends on the adjustment size and the closeness to real-time. Trading faster is preferred when larger adjustments are needed closer to real-time (Cramton et al., 2024).

13. Efficient trading

We envision five years of annual forwards (by hour, weekday-weekend), 12 months of monthly forwards (by hour and weekday-weekend), and 30 days of hourly forwards for each region and product type (premium, regular, and fast). This implies (3 product types) $\times$ (24 h/day) $\times$ [(2 day types) $\times$ (5 years +12 months) + 30] = 4608 products per region. With five hundred regions, this is 2.3 million products.

[6] The shading indicates price; the font color is to improve contrast. The numbers are simulated for illustrative purposes only. Detailed simulation of global communications markets is beyond the scope of this paper.

Yearly forward prices, New York, premium, weekday, $/GB

Hour	2033 (10)	2032 (9)	2031 (8)	2030 (7)	2029 (6)	2028 (5)	2027 (4)	2026 (3)	2025 (2)	2024 (1)
0	8.57	8.52	8.57	8.60	8.64	8.76	9.12	8.90	8.69	8.97
1	8.56	8.52	8.56	8.61	8.61	8.69	8.97	8.60	8.23	8.47
2	8.55	8.53	8.55	8.61	8.58	8.63	8.85	8.52	8.15	8.38
3	8.63	8.59	8.63	8.69	8.62	8.70	8.82	8.37	7.91	8.06
4	8.71	8.70	8.77	8.80	8.79	8.89	9.00	8.57	8.18	8.29
5	8.96	8.95	9.01	9.03	9.02	9.08	9.18	8.66	8.19	8.29
6	9.24	9.24	9.30	9.34	9.34	9.34	9.50	9.15	8.82	8.99
7	9.67	9.65	9.68	9.72	9.70	9.64	9.79	9.43	9.07	9.20
8	10.17	10.15	10.19	10.26	10.27	10.18	10.36	9.98	9.56	9.71
9	10.63	10.58	10.60	10.67	10.66	10.57	10.75	10.46	10.13	10.27
10	10.96	10.91	10.96	11.02	11.01	10.92	11.14	10.71	10.28	10.46
11	10.97	10.90	10.95	11.00	11.00	10.96	11.13	10.66	10.19	10.28
12	11.07	11.00	11.04	11.09	11.10	11.06	11.29	10.84	10.38	10.53
13	11.08	11.02	11.05	11.10	11.07	11.03	11.33	11.06	10.77	11.03
14	11.23	11.19	11.21	11.25	11.23	11.19	11.51	11.27	11.02	11.31
15	11.33	11.31	11.33	11.36	11.36	11.35	11.62	11.25	10.89	11.13
16	11.43	11.36	11.38	11.41	11.46	11.45	11.75	11.32	10.88	11.06
17	11.30	11.24	11.25	11.32	11.37	11.40	11.74	11.17	10.55	10.77
18	11.12	11.07	11.05	11.07	11.13	11.16	11.50	11.12	10.68	10.91
19	10.88	10.81	10.80	10.82	10.85	10.89	11.29	10.87	10.42	10.74
20	10.63	10.55	10.57	10.58	10.62	10.62	11.02	10.61	10.20	10.50
21	10.29	10.24	10.28	10.31	10.33	10.33	10.59	10.08	9.57	9.77
22	9.93	9.91	9.99	10.02	10.05	10.05	10.19	9.56	8.98	9.07
23	9.67	9.67	9.75	9.79	9.84	9.84	9.87	8.97	8.11	8.09

Price $/GB: 7.91 – 11.75

Fig. 5. Yearly forward prices, New York, premium, weekday, $/GB, 240 products per region.

Monthly forward prices, New York, premium, weekday, $/GB

Hour	Dec (12)	Nov (11)	Oct (10)	Sep (9)	Aug (8)	Jul (7)	Jun (6)	May (5)	Apr (4)	Mar (3)	Feb (2)	Jan (1)
0	8.53	8.49	8.51	8.45	8.36	8.15	8.20	8.06	7.63	7.99	8.41	7.90
1	8.51	8.47	8.51	8.44	8.40	8.22	8.32	8.14	7.74	8.30	8.93	8.40
2	8.56	8.49	8.56	8.50	8.42	8.29	8.40	8.24	7.82	8.35	8.90	8.37
3	8.69	8.60	8.64	8.57	8.47	8.37	8.49	8.34	7.95	8.38	8.76	8.24
4	8.92	8.82	8.87	8.81	8.69	8.63	8.73	8.62	8.24	8.62	8.90	8.43
5	9.12	9.04	9.09	9.05	8.93	8.84	8.90	8.80	8.48	9.06	9.57	9.21
6	9.36	9.30	9.38	9.35	9.25	9.15	9.15	9.09	8.84	9.65	10.42	10.16
7	9.73	9.70	9.74	9.68	9.58	9.47	9.46	9.32	9.09	9.93	10.75	10.48
8	10.21	10.21	10.25	10.20	10.10	9.95	9.92	9.78	9.57	10.41	11.24	11.01
9	10.60	10.60	10.67	10.65	10.57	10.45	10.44	10.28	10.04	10.97	11.83	11.59
10	10.91	10.92	10.97	10.97	10.87	10.78	10.76	10.80	10.51	11.63	12.63	12.34
11	10.93	10.94	10.97	10.99	10.90	10.86	10.84	10.86	10.58	11.44	12.14	11.88
12	11.07	11.07	11.10	11.09	10.98	10.92	10.86	10.97	10.72	11.57	12.29	12.07
13	11.11	11.08	11.12	11.08	10.98	10.91	10.85	10.89	10.78	11.58	12.30	12.20
14	11.27	11.23	11.32	11.26	11.16	11.04	10.98	10.99	10.82	11.71	12.56	12.40
15	11.39	11.35	11.43	11.40	11.29	11.16	11.05	10.98	10.85	11.75	12.64	12.60
16	11.43	11.41	11.45	11.47	11.33	11.21	11.07	10.94	10.66	11.48	12.25	12.13
17	11.29	11.29	11.32	11.34	11.19	11.08	10.99	10.86	10.53	11.51	12.42	12.21
18	11.09	11.08	11.10	11.08	10.94	10.86	10.81	10.67	10.33	11.43	12.43	12.12
19	10.89	10.89	10.93	10.89	10.79	10.70	10.65	10.59	10.23	11.23	12.18	11.84
20	10.64	10.61	10.68	10.65	10.55	10.43	10.34	10.26	10.03	10.91	11.75	11.58
21	10.27	10.26	10.34	10.34	10.24	10.14	10.05	10.07	9.88	10.65	11.38	11.28
22	9.88	9.89	9.96	9.97	9.85	9.71	9.62	9.63	9.45	10.20	10.92	10.81
23	9.60	9.62	9.71	9.73	9.60	9.47	9.43	9.47	9.29	10.16	10.96	10.80

Price $/GB: 7.63 – 12.64

Fig. 6. Monthly forward prices, New York, premium, weekday, $/GB, 288 products per region.

Hourly forward prices (odd days), New York, premium, $/GB

Hour	Days Ahead 29	27	25	23	21	19	17	15	13	11	9	7	5	3	1
0	8.55	8.57	8.58	3.17	8.55	8.59	3.24	8.68	8.64	8.68	2.49	8.76	8.06	2.39	7.64
1	8.54	8.55	8.56	3.12	8.52	8.56	3.20	8.62	8.59	8.64	2.37	8.74	8.20	2.26	7.87
2	8.56	8.58	8.58	3.08	8.51	8.56	3.16	8.61	8.57	8.60	2.43	8.77	8.43	2.40	8.24
3	8.62	8.65	8.66	3.08	8.57	8.58	3.09	8.62	8.59	8.61	2.54	8.78	8.34	2.66	8.09
4	8.74	8.78	8.82	3.26	8.73	8.72	3.18	8.77	8.70	8.72	2.62	8.79	8.30	2.65	8.02
5	8.94	8.98	9.04	3.51	8.96	8.92	3.41	8.99	8.92	8.99	2.62	9.04	8.53	2.46	8.24
6	9.26	9.27	9.32	3.76	9.26	9.19	3.75	9.27	9.19	9.30	3.06	9.36	8.78	2.84	8.34
7	9.67	9.67	9.70	4.23	9.66	9.59	4.30	9.70	9.60	9.69	3.61	9.77	9.36	3.26	9.07
8	10.16	10.16	10.17	4.72	10.15	10.11	4.78	10.19	10.06	10.11	4.03	10.12	9.64	3.64	9.30
9	10.59	10.58	10.60	5.18	10.60	10.59	5.24	10.68	10.52	10.56	4.42	10.48	9.92	3.90	9.57
10	10.88	10.89	10.92	5.44	10.92	10.92	5.49	11.01	10.88	10.89	4.70	10.72	10.16	4.22	9.83
11	10.91	10.96	10.99	5.46	10.95	10.95	5.54	11.04	10.97	10.97	4.78	10.81	10.23	4.26	9.95
12	11.02	11.08	11.11	5.58	11.05	11.01	5.65	11.07	11.04	11.02	4.96	10.89	10.57	4.43	10.56
13	11.06	11.12	11.14	5.57	11.03	11.00	5.63	11.06	11.05	10.96	5.01	10.81	10.56	4.61	10.58
14	11.24	11.27	11.26	5.70	11.20	11.16	5.76	11.21	11.18	11.16	5.39	11.14	10.83	5.19	10.72
15	11.38	11.36	11.31	5.76	11.27	11.26	5.86	11.31	11.26	11.25	5.67	11.32	10.86	5.72	10.53
16	11.41	11.38	11.33	5.79	11.34	11.32	5.86	11.38	11.33	11.40	5.64	11.53	11.01	5.69	10.65
17	11.25	11.27	11.23	5.66	11.21	11.18	5.70	11.26	11.22	11.32	5.56	11.38	10.81	5.60	10.40
18	11.04	11.09	11.08	5.49	11.07	11.01	5.48	11.07	11.00	11.14	5.06	11.20	10.47	4.94	9.91
19	10.83	10.89	10.92	5.39	10.88	10.78	5.40	10.84	10.77	10.85	4.66	10.88	10.08	4.31	9.53
20	10.58	10.62	10.64	5.18	10.59	10.50	5.20	10.57	10.50	10.58	4.36	10.60	10.01	4.00	9.68
21	10.28	10.29	10.28	4.81	10.25	10.16	4.86	10.27	10.20	10.28	4.26	10.24	9.72	3.98	9.49
22	9.94	9.91	9.87	4.35	9.84	9.79	4.38	9.93	9.87	9.90	4.15	9.81	9.35	3.94	9.04
23	9.67	9.64	9.59	4.05	9.58	9.55	4.06	9.68	9.64	9.63	4.10	9.60	8.99	4.06	8.47

Price $/GB: 2.26 – 11.53

Fig. 7. Hourly forward prices, New York, premium, $/GB, 720 products per region.

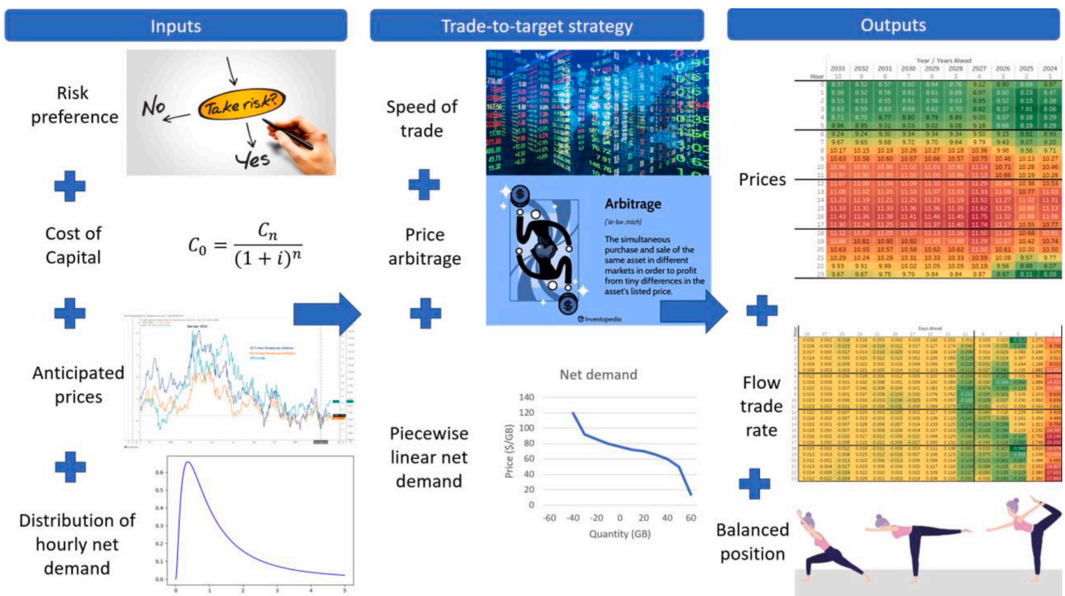

Fig. 8. A participant's inputs imply a trade-to-target strategy and resulting outputs.

We can summarize the key features of the bidding language in two theorems from Budish et al. (2023) and an immediate corollary. The mathematics below borrows freely from Budish et al. (2023).

Let $V_i(x_i)$ denote the dollar utility of order i from a trade rate of $x_i = D_i(p_i)$ in portfolio units per hour, where flow portfolio demand $D_i(p_i)$ is given by equation (1):

$$D_i\left(p_i \middle| \mathbf{w}_i, q_i, p_i^L, p_i^H\right) := q_i trunc\left(\frac{p_i^H - p_i}{p_i^H - p_i^L}\right) \text{ where } trunc(z) := \begin{cases} 1 \text{ for } z \geq 1 \\ z \text{ for } 0 \leq z < 1 \\ 0 \text{ for } z \leq 0 \end{cases}. \tag{1}$$

To find $V_i(x_i)$, we first define the marginal utility function $M_i(x_i)$ as the inverse demand curve, $p_i = M_i(x_i)$. The inverse demand curve maps order i's trade rate $x_i \in [0, q_i]$ into prices $p_i \in [p_i^L, p_i^H]$.[7] Rearranging equation (1), we have:

$$M_i(x_i) := p_i^H - \frac{p_i^H - p_i^L}{q_i} x_i \quad \text{for } x_i \in [0, q_i]. \tag{2}$$

The value of $M_i(x_i)$ measures marginal as-bid flow value in dollars per portfolio unit. Utility $V_i(x_i)$, as a function of the trade rate x_i, is defined as the integral of the marginal utility function over the interval $[0, x_i]$:

$$V_i(x_i) := \int_0^{x_i} M_i(u)\, du. \tag{3}$$

Since the marginal value is linear in x_i, the total value is quadratic and strictly concave in x_i:

$$V_i(x_i) = p_i^H x_i - \frac{p_i^H - p_i^L}{2q_i} x_i^2. \tag{4}$$

We assume $V_i(x_i)$ as defined for all $x_i \in \mathbb{R}$, with order specifications imposing the constraint $x_i \in [0, q_i]$.

Our problem of finding market-clearing prices is formulated as two optimization problems: a primal problem of finding quantities that maximize as-bid dollar value and a dual problem of finding prices that minimize the cost of non-clearing prices. The first-order conditions for the optimality of these two problems imply market-clearing prices and quantities.

The market operator, acting analogously to a social planner, chooses a vector of trade rates for all orders $\boldsymbol{x} = (x_1, \ldots, x_I)$ to maximize aggregate value, defined as the sum of pseudo-utility functions across orders,

$$V(\boldsymbol{x}) := \sum_{i=1}^{I} V_i(x_i) \quad \text{for } \boldsymbol{x} \in \mathbb{R}^I, \tag{5}$$

subject to market-clearing constraints and trade-rate constraints:

$$\max_{\boldsymbol{x}} V(\boldsymbol{x}) \quad \text{subject to} \begin{cases} \sum_{i=0}^{I} x_i \boldsymbol{w}_i = \boldsymbol{0} & \text{(market-)} \\ x_i \in [0, q_i] \text{ for all } i & \text{(trade-).} \end{cases} \tag{6}$$

The objective function $V(\boldsymbol{x})$ is concave because it is a sum of concave functions.

Indeed, this is a quadratic program since the objective function is quadratic and the constraints are linear. To make this quadratic structure apparent using matrix and vector notation, let $\boldsymbol{W}$ denote the $N \times I$ matrix whose i th column is $\boldsymbol{w}_i$. Let $\boldsymbol{p}^H$ denote the column vector whose i th element is p_i^H. Let $\boldsymbol{D}$ denote the $I \times I$ positive definite diagonal matrix whose i th diagonal element is $(p_i^H - p_i^L)/\, q_i$. Then, the problem in equation (6) may be written compactly as

$$\max_{\boldsymbol{x}} \left[\boldsymbol{x}^T \boldsymbol{p}^H - \frac{1}{2}\, \boldsymbol{x}^T \boldsymbol{D} \boldsymbol{x}\right] \quad \text{subject to} \quad \boldsymbol{W}\boldsymbol{x} = \boldsymbol{0} \quad \text{and} \quad \boldsymbol{0} \le \boldsymbol{x} \le \boldsymbol{q}. \tag{7}$$

We first show that quantities that maximize aggregate utility exist. Then, we show that market-clearing prices exist by examining the dual problem of the utility maximization problem. We then show that there is a unique mapping of orders into prices and quantities. Uniqueness of prices and quantities is important for transparency. These are standard results of convex optimization (Bertsekas, 2009), derived from strict convexity and continuity. Our presentation follows Budish et al. (2023).

Theorem 1. Existence and Uniqueness of Optimal Quantities. A unique vector of trade rates **x** exists, which solves the maximization problem in equation (7).

To prove that market-clearing prices exist, we exploit the duality between the problems of finding optimal prices and quantities. For this, we define a Lagrangian function of the vector of trade rates $\boldsymbol{x}$ with three constraints: (1) the market clears ($W\boldsymbol{x} = \boldsymbol{0}$); (2) the trade rates are greater than or equal to zero ($\boldsymbol{x} \ge \boldsymbol{0}$); (3) the trade rates are less than or equal to their maxima ($\boldsymbol{x} \le \boldsymbol{q}$). In vector notation, the Lagrangian is defined by

$$L(\boldsymbol{x}, \pi, \lambda, \mu) := \boldsymbol{x}^T \boldsymbol{p}^H - \frac{1}{2}\, \boldsymbol{x}^T \boldsymbol{D}\boldsymbol{x} - \pi^T \boldsymbol{W}\boldsymbol{x} + \mu^T \boldsymbol{x} + \lambda^T (\boldsymbol{q} - \boldsymbol{x}). \tag{8}$$

Since the multipliers associated with the market-clearing equality constraints have the economic interpretation of market prices for assets, we use the notation $\pi = (\pi_1, \ldots, \pi_N)^\top$ for these multipliers. Two vectors of multipliers, $\mu = (u_1, \ldots, \mu_I)^\top$ and $\lambda = (\lambda_1, \ldots, \lambda_I)^\top$, are associated with inequality constraints on trade rates.

The dual problem associated with the primal problem of maximizing aggregate utility in equation (7) is then defined by

[7] For trade rates in the interval $(0, q_i)$, the fact that the order chooses an interior trade rate tells us that the order's as-bid marginal utility is equal to the corresponding price in the interval (p_i^L, p_i^H). The same logic extends to the boundary points 0 and q_i, corresponding respectively to prices p_i^H and p_i^L, by assuming as-bid utility is continuous.

$$\widehat{G}(\pi,\lambda,\mu) : = \max_{x} L(\boldsymbol{x},\pi,\lambda,\mu) \quad \text{for} \quad \pi \in \mathbb{R}^N, \mu \geq \mathbf{0}, \lambda \geq \mathbf{0}. \tag{9}$$

The dual problem is a minimization problem with infimum g defined by

$$g : = \inf_{\pi,\lambda,\mu} \widehat{G}(\pi,\lambda,\mu) \quad \text{subject to} \quad \pi \in \mathbb{R}^N, \mu \geq \mathbf{0}, \lambda \geq \mathbf{0}. \tag{10}$$

The dual problem in equation (10) is formulated as an infimum rather than a minimum because we have not yet shown that there exists a solution (π, λ, μ) that attains the infimum.

Theorem 2. *Existence of market-clearing. There exists at least one optimal solution* (π,λ,μ) *to the dual problem in* equation (10). *The solutions* $\boldsymbol{x}$ *and* (π,λ,μ) *are a primal-dual pair which satisfies the strict duality relationship*

$$g = V(\boldsymbol{x}). \tag{11}$$

Theorem 2 does not guarantee that market-clearing prices are unique. The set of market-clearing prices is convex and may be unbounded. A trivial example occurs when all orders are buy orders for individual assets, and there are no sell orders. Then, any sufficiently high price clears the market with zero trade. There may also be cases where the market-clearing price is not unique even when trade occurs. A trivial example occurs when there is one buy order and one sell order for the same asset (or portfolio) with the same maximum rate, and the buyer's lower limit price exceeds the absolute value of the seller's lower limit price. In this case, there is an interval of prices where both orders are fully executable. However, a natural tie-breaking rule makes the prices unique.

Closest-to-prior-prices rule. If more than one price vector supports the optimal quantity vector, select the price vector closest to the prior price vector.

Corollary 1. *Uniqueness of quantities and prices. Prices and quantities are unique with the closest-to-prior-prices rule.*

Proof. *The set of prices that support the unique optimal quantities is convex. The closest point in a convex set to a point is unique. End proof.*

The closest-to-prior-prices tie-breaking rule is especially appropriate in our frequent batch auction setting, in which prices evolve slowly from the gradual trade of persistent orders.

These unique prices and quantities can be found quickly. Flow trading involves the solution of the following optimization program:

$$\min_{x} \quad \frac{1}{2}x^\top Dx - p^\top x \quad \text{s.t.} \quad a \leq x \leq b \quad \text{and} \quad Wx = 0,$$

where D is a non-negative, diagonal matrix. To exploit the near-separability of the problem, we employ the alternating direction method of multipliers (ADMM) (Boyd et al., 2011). This technique solves an optimization problem of the form

$$\min_{x,z} \quad f(x) + g(z) \quad \text{s.t.} \quad Ax + Bz = c.$$

We define an indicator function $C(b) = 0$ if b is true and ∞ otherwise, i.e., $C(a \leq x \leq b)$ and $C(Wx = 0)$ will be used to enforce our problem constraints. We choose

$$f(x) = \frac{1}{2}x^\top Dx - p^\top x + C(a \leq x \leq b),$$

$$g(z) = C((1^\top \otimes I)z = 0),$$

and

$$A = \sum_i (e_i e_i^\top) \otimes (We_i), B = -I, c = 0,$$

where $\otimes$ denotes the Kronecker product. This splits the minimization across two sets of variables: x, which correspond to rates of execution of each order, and $z = (z_1, z_2, \ldots)$, which are the trade rates each fulfilled order imposes across the space of products, i.e. $(We_i)x_i = z_i$ for each order i.

ADMM proceeds by formulating the augmented Lagrangian, then repeatedly minimizing it via a Gauss-Seidel pass on the primal variables (x,z), followed by a dual ascent on y. When substituted into the ADMM framework, the splitting scheme yields a compelling algorithm. First, it is straightforward to show that the dual variable $y = 1 \otimes \pi$, where π are the shadow prices. Second, the two subproblems needing to be solved as part of the Gauss-Seidel pass are trivial.

The first subproblem, necessary for the x-update, takes the form

$$\min_{x} \quad \frac{1}{2}x^\top (D + \rho I)x - r^\top x \quad \text{s.t.} \quad a \leq x \leq b,$$

for some r that varies per iteration. Being fully separable, we can write the solution explicitly:

$$x_i = \max\left(a_i, \min\left(b_i, r_i / (D_{ii} + \rho)\right)\right).$$

The second subproblem, necessary for the z-update, takes the form

$$\min_{\{z_k\}} \sum_k \frac{1}{2} z_k^\top z_k - c_k^\top z_k \quad \text{s.t.} \quad \sum_k z_k = 0,$$

where the $\{c_k\}$ vary per iteration. This can be solved analytically using elementary calculus and evaluated by simply averaging the $\{c_k\}$.

Both subproblems efficiently scale to arbitrarily large problem sizes and are easily parallelized on a CPU or GPU. Research is ongoing to fine-tune the implementation's penalty and over-relaxation parameters.

We have built a prototype platform to implement forward markets in many domains, such as energy, communications, and transportation. The basic architecture is depicted in Fig. 9. The core infrastructure is a forward market system and a low-level flow trading system that performs optimization.

Details of flow trading are shown in Fig. 10. It consists of an application programming interface, a database, and an optimization engine.

Although it may be hard to imagine trading so many products, the flow trading technology makes this easy by exploiting the power of convex optimization. From a computational complexity viewpoint or a user experience viewpoint, there is little difference between a hundred, a thousand, or a million products. Bid entry and optimization are readily managed with information technology. Budish et al. (2023) demonstrate how computation times vary with the number of orders and assets (products); see Fig. 11. The computation to find unique prices and quantities can be done on a single server in about one hundred seconds, allowing clearing every hour. The Appendix includes sample code for a simple flow trading implementation.

14. Simplified participation in a complex market

Thanks to a simple and effective method of preference expression, participation is straightforward, even though the market is solving for a complex set of dynamic demand and supply variables. A participant's strategy depends on three essential inputs: risk attitude, capital cost, and expected hourly net demand.

Standard financial modeling provides the simplest way to represent risk attitude and capital cost. Two scalar parameters can define the participant's utility function. Capital cost is the participant's discount rate or time value of money. Risk attitude determines the concavity of the utility function. Assuming constant absolute risk aversion, a risk-neutral participant would have a risk parameter of zero, implying linear utility. Risk-averse participants have a risk parameter greater than zero. Larger risk parameters indicate greater risk aversion.

The final input is the participant's hourly expected net demand. This is easy for a pure financial participant; net demand is zero for all hours. For others, it is a complex technical calculation that requires good knowledge of customers for buyers and portfolio for sellers. Hybrid participants who own a portfolio of capacity tokens and serve consumers must estimate their anticipated supply and demand. However, participants in any communications wholesale market need to estimate net demand to decide how to participate. This difficult input is necessary regardless of the market's design (Cramton et al., 2024).

A trading tool is developed to help market participants understand how the forward market works and how to participate easily. The core of the tool is straightforward. The user specifies the risk attitude and capital cost parameters. The user then uploads her expected net demand as a.csv file. For an arbitrageur, this is a matrix of zeros; no upload is needed in this case since the zero matrix is

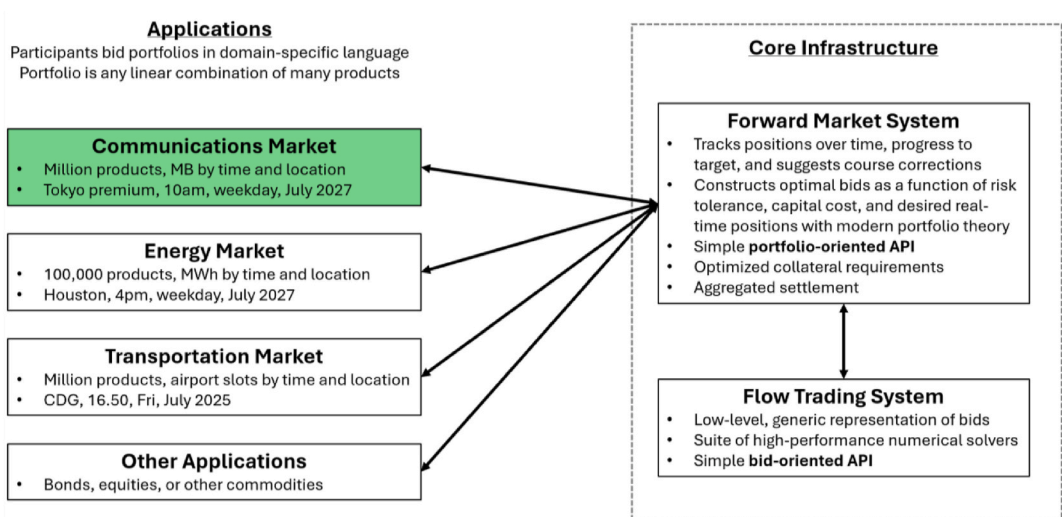

Fig. 9. Forward market architecture.

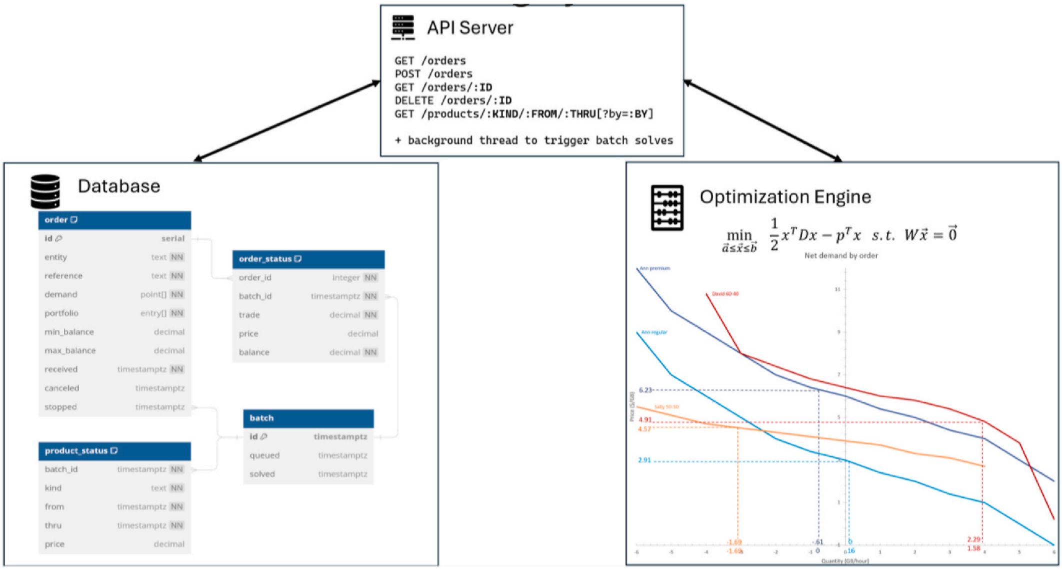

Fig. 10. The flow trading system.

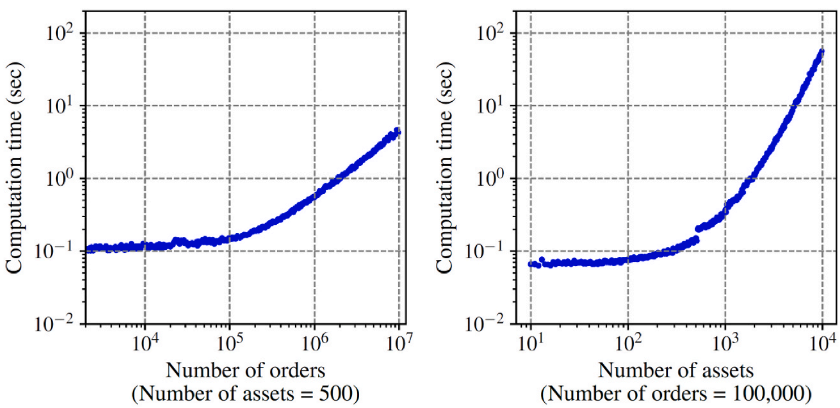

Fig. 11. Computation time for challenging cases by number of orders and assets (products) (Budish et al., 2023); (the slight discontinuity in the right panel has to do with the cache size of the CPU).

the default. The arbitrageur's target position is zero for all products. For a buyer and seller (or hybrid), the trader's expected net demand (demand minus supply) defines the target position in each hour. The target is the participant's expected net demand in each hour multiplied by the trader's target percentage. The target percentage increases linearly from 0 percent to 100 percent from five years ahead until real-time.

As time passes, uncertainty resolves, and the participant adjusts its target strategy. The adjustments are modest. This is a simulation of a participant of moderate size (about 240 GB/h in each region).

Fig. 12 shows how the target can be reached with a flow trade rate, assuming a flow trade rate as (communications adjustment)/(8 × days ahead). Flow trade rates are small if the days ahead are large. There are many hours to trade when we are far from real-time. This is why the flow trade rates are so small many days from real-time. Even close to real-time, the required flows are only a handful of gigabytes, which is small for a 240 GB/h service provider. The quantity traded is never zero but always small, reducing risk and adverse price impact.

The output lets the user visualize outcomes and how outcomes vary with variations in specified risk attitude and capital cost. It also helps users determine the incremental gain or loss from adding customers (increasing demand). This incremental calculation is essential in pricing and investment decisions.

Flow trade rate in straightforward strategy (GB/hour)

Hour \ Days Ahead	29	27	25	23	21	19	17	15	13	11	9	7	5	3	1
0	0.031	0.002	-0.018	0.016	-0.003	0.002	0.033	0.142	0.201	0.001	0.020	0.023	-0.312	1.271	10.510
1	0.026	-0.003	-0.023	0.006	-0.019	-0.012	0.017	0.127	0.174	-0.040	0.009	-0.035	-0.156	1.139	8.798
2	0.017	0.000	-0.017	0.013	-0.019	-0.020	-0.002	0.108	0.119	-0.099	0.011	-0.023	0.089	0.280	3.070
3	0.014	0.002	-0.005	0.020	-0.012	-0.010	0.001	0.106	0.086	-0.119	-0.005	-0.024	0.307	0.438	3.511
4	0.009	0.000	-0.009	0.022	0.000	0.016	0.006	0.099	0.067	-0.128	-0.031	-0.103	0.244	1.186	8.013
5	0.009	0.003	-0.003	0.021	-0.005	0.014	0.010	0.103	0.085	-0.112	-0.027	-0.141	0.135	1.880	11.952
6	0.012	0.003	-0.006	0.028	-0.006	0.012	0.005	0.100	0.084	-0.104	-0.023	-0.162	0.005	1.960	13.054
7	0.019	0.009	0.001	0.032	-0.008	0.001	0.009	0.100	0.089	-0.126	-0.030	-0.199	-0.160	1.886	14.221
8	0.022	0.011	0.007	0.040	-0.009	-0.009	-0.001	0.083	0.053	-0.189	-0.074	-0.163	-0.124	1.339	11.296
9	0.023	0.009	0.001	0.037	-0.009	-0.024	-0.010	0.076	0.052	-0.210	-0.095	-0.101	0.105	1.203	8.633
10	0.023	0.003	-0.006	0.033	-0.013	-0.036	-0.025	0.076	0.057	-0.195	-0.109	-0.037	0.230	0.992	6.839
11	0.023	0.002	-0.011	0.028	-0.008	-0.031	-0.008	0.094	0.079	-0.170	-0.081	-0.002	0.334	0.543	3.842
12	0.022	0.003	-0.007	0.029	-0.003	-0.020	0.002	0.117	0.104	-0.149	-0.043	0.018	0.229	0.499	4.418
13	0.019	0.002	-0.007	0.023	-0.003	-0.003	0.029	0.139	0.119	-0.141	-0.073	-0.069	0.181	0.454	4.499
14	0.015	-0.001	-0.007	0.028	-0.006	-0.007	0.014	0.133	0.120	-0.148	-0.126	-0.156	0.060	1.311	9.754
15	0.015	0.000	-0.007	0.022	-0.009	-0.009	-0.004	0.107	0.119	-0.135	-0.133	-0.156	0.133	2.231	14.747
16	0.016	-0.002	-0.010	0.017	-0.012	-0.009	-0.010	0.099	0.129	-0.096	-0.091	-0.159	-0.147	2.755	19.144
17	0.016	0.000	-0.008	0.015	-0.011	-0.009	-0.017	0.095	0.131	-0.100	-0.045	-0.118	-0.182	2.748	20.250
18	0.015	-0.002	-0.012	0.019	-0.011	-0.009	-0.003	0.111	0.138	-0.082	-0.033	-0.107	-0.346	1.610	14.574
19	0.013	0.003	-0.008	0.025	-0.012	-0.016	-0.007	0.106	0.123	-0.113	-0.075	-0.110	-0.201	1.039	10.604
20	0.011	-0.001	-0.013	0.022	-0.006	-0.013	0.006	0.099	0.114	-0.129	-0.121	-0.161	-0.167	0.486	6.688
21	0.013	-0.004	-0.017	0.023	0.006	-0.004	0.020	0.117	0.116	-0.134	-0.088	-0.119	-0.050	2.052	14.807
22	0.014	-0.008	-0.024	0.022	0.010	0.002	0.022	0.109	0.118	-0.117	-0.050	-0.128	-0.119	2.380	17.653
23	0.012	-0.012	-0.024	0.029	0.011	-0.003	0.014	0.106	0.115	-0.091	-0.006	-0.143	-0.153	2.360	17.893

Flow trade rate (G..
-0.346 20.250

Fig. 12. Flow trade rate at prior communications price (GB/hour).

15. Liquidity, counterparty risk, flexibility, and competition

15.1. Liquidity

Traditional markets manage liquidity by limiting the number of products. For example, wheat trading involves many grades and classifications, which vary by country and the organization responsible for grading. The United States Department of Agriculture categorizes wheat into eight classes based on kernel hardness, color, and planting season. Within these classes, wheat is further graded on a scale from 1 through 5 based on additional attributes like test weight, defects, and moisture content. There are forty wheat products traded in the US.

As explained above, modern markets like the open-access market can trade products with much richer granularity. Liquidity is managed by allowing near-perfect substitution among products that are near-perfect substitutes. Through gradual adjustment of portfolios, the global capacity tokens become fungible network capacity for communications at particular times and locations.

The forward market has high transparency, robust pricing, and low transaction costs, which favors liquidity. The forward market has three further advantages. First, preferences are convex. Market participants enter piecewise linear net demand curves, which yield a quadratic objective in the clearing optimization. Second, because the forward market is conducted well before the real-time market, the market participants have time to adjust positions as uncertainty resolves. Third, the frequent batch auction approach allows participants to make thousands of minor adjustments over months and years. Slow trading enables participants to minimize adverse price movements, improving the market's competitiveness and increasing liquidity.

15.2. Counterparty risk

Efficient and transparent forward trade reduces counterparty risk and lowers costs. Vibrant forward trade puts market participants in more balanced positions, reducing risk and market power and thereby reducing system cost.

Electricity markets provide a vivid example of the benefit of balanced positions in reducing counterparty risk. Consider the costly defaults in electricity markets over the last twenty years. In the 2000–2001 California electricity crisis, the utilities entered a long scarcity period caused by drought (low hydro production) with a large short position (Borenstein, 2002). The utilities required rescue by the state, costing about $40 billion (California State Auditor, 2001). In the February 2021 Texas crisis, the market participants were in more balanced positions, and defaults were rare despite a real-time value of electricity of over $50 billion in four days (Cramton, 2022). In Britain's crisis of 2021–2022, poorly hedged suppliers defaulted, costing consumers more than $10 billion (Waddams, 2023).

In the forward market, imbalanced positions are known, and the associated risk is priced and mitigated through higher collateral. Overall system risk is reduced.

15.3. Open access motivates flexibility

Efficient prices reward those providing flexibility. Market participants can easily see and enjoy the value of flexibility. Transparent and efficient prices will motivate demand-side innovation essential to consumer engagement (Cramton et al., 2024).

15.4. Open access is pro-competitive

A few concentrated firms, static pricing and services, and a need for more innovation and competition are hallmarks of today's communications markets. With an open, nondiscriminatory wholesale market, enhanced through granular real-time and forward market pricing mechanisms that consider time and place, more fungible network capacity is made available at more affordable prices. This will improve the ability of smaller and more innovative communications operators to be competitive in their respective markets. As noted above, it will also encourage operators to be more flexible with their capacity portfolios, prompting them to sell excess capacity when it is not needed rather than strand it behind a walled garden.

16. Summary and conclusions

This paper develops an open-access market to manage network congestion and optimize a network's use and value, building on recent advances in wholesale electricity markets (Cramton et al., 2024). The wholesale market includes real-time and forward markets. The initial conceptual application is intersatellite communication networks—optical (laser-beamed) mesh networks in space.

This physical real-time market with priority pricing ensures a balance between offered supply and bid demand at each time and location. The product, represented by a crypto token, is gigabytes of three communication types—premium, regular, and fast—in a 1-h time window at a location, say, New York City, 9–10am on August 14, 2026. The scarce capacity is used by those who value it the most at a price that balances supply and demand. The real-time market is the foundation for the financial forward market. The forward market enables market participants to take capacity positions in advance of real-time, consistent with their anticipated real-time needs. Participants manage risk and profit through gradual trade as uncertainties resolve. Participants can efficiently convert global communication rights into their realized communication needs at each time and place. An independent market operator conducts a transparent market.

The open-access market operates without friction using flow trading (Budish et al., 2023). Participants bid persistent piecewise-linear downward-sloping net demand curves for portfolios of products. The market operator clears the market every hour, finding unique prices and quantities that maximize as-bid social welfare. Prices, aggregate quantities, and the slope of the aggregate net demand are made public through a blockchain platform. The market operator observes positions, enabling it to optimize collateral requirements to minimize default risk.

Participants employ trade-to-target strategies with only a handful of portfolio orders, enabling the efficient trade of millions of interrelated time-and-location-specific products. In each hour, the participant has a current position—the quantity held of each product—and a target position—the desired portfolio. The trade-to-target strategy specifies the rate at which the participant moves from its current position toward its target. To best manage risk and avoid adverse price impact, participants trade gradually, updating their target as circumstances change. Participants' strategies depend on their communication needs, risk tolerance, capital cost, and market fundamentals. Fundamentals are richly conveyed in the market operator's hourly clearing reports. The market operator provides tools for participants to translate their preferences into an effective strategy. Despite the complexity of the market, participation is easy.

Market participants also have an opportunity to take ownership and governance roles in the market by acquiring and trading governance tokens, which are also bought and sold through the market's blockchain platform. These tokens represent a bundle of rights akin to rights attached to shares held by shareholders in a private or public company. Holders of governance tokens are eligible to receive allocations of market revenues and to vote in elections for representatives to the market's governing body.

The ideas presented here are familiar. The most important points have been well-understood for decades if not centuries. The inefficiencies created by imbalanced ownership appear in Myerson and Satterthwaite (1981), Cramton et al. (1987), and Ausubel et al. (2014) and are empirically documented in many studies (Borenstein, 2002; Wolak, 2003; Wolfram, 1999). Vickrey (1961) pricing can mitigate market power but only with non-anonymous, discriminatory prices that seem unfair to many and are anticompetitive in favoring larger parties. More frequent trade provides a better means to mitigate market power (Black, 1971; Coase, 1972; Kyle, 1985; Vayanos, 1999), especially when dynamic trade is natural to manage risk as circumstances change. The form of trade matters. Frequent batch auctions can eliminate an arms race for speed (Budish et al., 2015) and have been implemented in electricity markets (Cramton & Ockenfels, 2024), especially when combined with flow trading (Budish et al., 2023), which brings an effective language of preference expression. Participants can adopt simple trade-to-target strategies, allowing flexible risk management and providing efficient and transparent price signals. These prices summarize the essential information for efficient investment (Cramton et al., 2024).

The ideas apply to any commodity, especially those with time and location elements. In this paper, we have elaborated on trade in optical mesh networks in orbit for intersatellite communications. Considering other markets, the most obvious is mobile communications. Price is a more efficient instrument for managing congestion than rationing with dropped calls or throttling (Cramton & Doyle, 2017).

Infrequent spectrum auctions are used to assign mobile communications spectrum today. These auctions could be replaced by a much more flexible assignment of spectrum in time and place. We need communication technologies that are sufficiently flexible. Then, spectrum can be a commodity traded in real time to balance supply and demand, supported by a forward spectrum market.

Doing so would address the market power and competition issues of today's oligopoly model. Rather than continued consolidation, mobile communications could shift to an open-access model of vibrant competition like the internet enabling Eli Noam's vision for open access in spectrum-based communications markets (Noam, 1998). Communications is the natural product, since that is the commodity buyers consume. However, as devices become more flexible, it may be possible to commoditize spectrum over short time intervals. Then, there could be multiple commodities: real-time markets for communications and spectrum, with spectrum serving as an input market for the communications end product.

Various organizational structures are possible for the communications and spectrum markets. The communications market could be conducted by a neutral operator organized by one or more network owners. The spectrum market, as an essential input of communications, could be conducted by a similar neutral market operator; however, the spectrum market would be more tightly regulated by the communications regulator. Wholesale electricity markets provide an example. The electricity regulator designates an independent system operator to conduct the open access electricity market to provide reliable electricity at least cost. This structure has worked well in restructured electricity markets (Cramton, 2017). Natural gas is an input to electricity production, while spectrum is an input to communications. In the future, real-time and forward spectrum markets may replace today's spectrum auctions, fulfilling Noam's (1998) vision that spectrum auctions are tomorrow's anachronism.

We have stressed the benefits of open access. But what about the costs? Perhaps the costs may dominate the significant benefits. The answer is that historically, serious costs stood in the way of open access. Information technology advances have eliminated these costs. The open access benefits can be enjoyed today without cost.

An open-access market for global communications offers an early entry into the brave new world of efficient pricing of essential commodities. Transparent and efficient pricing benefits market participants by maximizing the value of the scarce communications capacity. The prices also provide crucial information for efficient investment and operation of network resources.

CRediT authorship contribution statement

Peter Cramton: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Erik Bohlin:** Writing – review & editing. **Simon Brandkamp:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Jason Dark:** Writing – review & editing, Writing – original draft, Software, Methodology, Investigation, Formal analysis. **Darrell Hoy:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Albert S. Kyle:** Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **David Malec:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Axel Ockenfels:** Writing – review & editing, Supervision, Resources, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Chris Wilkens:** Writing – review & editing, Visualization, Methodology, Conceptualization.

Acknowledgements

Support by Rivada Networks and the German Science Foundation through Germany's Excellence Strategy (EXC 2126/1 390838866) is gratefully acknowledged.

Appendix

Sample source code that implements flow trading

For a zip file with sample source code, see https://cramton.umd.edu/communications and click on Sample Code. The sample code does not include the advanced tricks, decompositions, or preconditioning performed in a production-grade implementation. It is intended to reflect the mathematical logic of flow trading and help validate or benchmark more advanced deployments.

If flow trading is so great, why has it not been adopted in financial markets?

Financial markets suffer from the same limitations of the stakeholder process as communications markets. The most dominant stakeholders lobby the regulators to adopt market rules that favor them. For example, high-frequency traders dominate the technical committee advising the Commodity Futures Trading Commission. It is no wonder that the CFTC is slow to adopt reforms that limit the profits enjoyed by traders with speed advantages, especially since the other influential stakeholders, the exchanges, make most of their money selling tools—data and collocation services—to high-frequency traders. There is little tendency for the market to adopt efficiency-enhancing reforms. Regulators are risk-averse and easily scared that a reform may have unintended consequences. Such is the tyranny of the status quo (Budish et al., 2015).

Indeed, the market design challenges in financial markets are worse than in communications markets. Important financial stakeholders were entrenched when information technology made the reforms discussed here possible. Thus, these stakeholders provided immediate resistance to change. A norm of transparency did not exist and, indeed, was prevented by legislation from the early 1900s that prevents the disclosure of bidding in Treasury markets—even after one hundred years.

Do you need to worry about a short squeeze as in other forward commodity markets?

The forward market settles against the real-time price. All forward products are financial derivatives of the physical real-time products.

However, as in any forward market, a short squeeze is possible. The squeeze would take place in the real-time market. A participant takes a significant imbalanced position in the forward market, causing others to take imbalanced forward positions and then squeeze them in real-time. A dominant supplier has a comparative advantage in executing the squeeze. A supplier buys a large quantity forward, leaving others short, then offers supply at high prices in real-time and strategically withholds in real-time. During periods of scarcity, enhanced market power improves the effectiveness of such strategies.

The forward market mitigates this possibility through transparency of positions. The market operator and market monitor would observe the imbalanced position, prompting action. Moreover, the single-price auction makes a squeeze prohibitively expensive.

Recall Salomon Brothers' famous squeeze in the US Treasury markets in 1990–1991. To be successful, Salomon Brothers needed to hold a considerable position. They acquired majority shares in some Treasury auctions. Although illegal, winning a majority was possible because of the pay-*as*-bid pricing and large price-tick size at the time. Salomon Brothers could acquire most of the issue and squeeze the short dealers in the subsequent market by bidding at one tick above the obvious clearing price. Acquiring such a significant stake would be prohibitively expensive with single pricing, which we have here.

Hundreds of market participants exist in the forward communications market. The participants include natural buyers, natural sellers, and arbitragers. Natural buyers and sellers also function as arbitragers—the arbitrage behavior results in price convergence. In electricity markets, the forward price equals the expected real-time price plus a small risk premium of less than two percent (Jha & Wolak, 2023).

The market is highly competitive. Therefore, the scope for strategic bidding is limited. Flow trading further mitigates incentives for strategic bidding by incentivizing participants to seek balanced positions to manage risk and limit collateral. With balanced positions, there is no incentive to distort bids.

Market power only arises close to real-time. Then, market participants can take actions that may result in more significant and favorable price impacts because other participants will not have time to take corrective measures to mitigate this behavior.

Example of preference expression and market clearing

To fix ideas, consider three market participants (Ann, David, and Sally), two locations (New York and Tokyo), and two times (today and tomorrow). Our participants submit bids in today's forward market to hedge tomorrow's prices. Deviations from today's position will be realized tomorrow and settled at tomorrow's prices.

Ann is an arbitrageur. She participates in the market to exploit her expert understanding of prices. Her strategy is classic: buy low and sell high; do not drift far from a zero position.

Sally is a US-based communications service provider with a portfolio of capacity tokens that she sometimes deploys for domestic communications operations. Her portfolio exceeds her needs and often sells capacity into the open-access market. Sally pursues a trade-to-target strategy designed to maximize profit and limit risk.

David is a wireless reseller. He has a portfolio of consumers he is obligated to serve but never wants to hold excess capacity. He participates in the market to maximize profit and limit risk.

Although Sally is a seller and David is a demander, both recognize that it is helpful, like Ann, to participate in buying and selling depending on prices and other circumstances. Thus, all market participants express net demand curves that involve selling or buying depending on price. Participants express quantity as a flow, the rate of trade over a 1-h window (GB/hour).

First, suppose there is a single forward product, tomorrow's premium capacity. Each demand curve is expressed as a vector of quantity-price pairs, as in Table A1.

Table A1
Piecewise linear net demands (GB/hour) as a function of price ($/GB)

Quantity GB/hour	Price ($/GB)		
	Ann	Sally	David
−6	12	7	
−5	10	6.6	
−4	9	6.2	12
−3	8	6	9.2
−2	7	5.8	8.6
−1	6.4	5.6	8
0	6	5.4	7.6
1	5.4	5.2	7.2
2	5	4.8	7
3	4.4	4.6	6.6
4	4	4.2	6
5	3		5
6	2		1.4

Figure A1 shows the net demand curve for each participant.

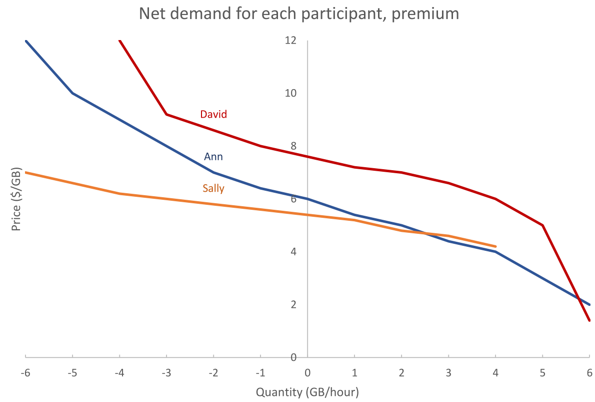

Fig. A1. Net demand curve for each participant.

Ann expects tomorrow's premium demand to be about $6/GB. She wants to buy when the forward price is less than $6 and sell when it is higher than $6. To protect herself from going too long or short, Ann bids a net demand curve that becomes steeper when the absolute value of the quantity is larger. This shape is a risk management element in all three curves: convex for negative amounts (buying) and concave for positive quantities (selling). It also mitigates adverse selection and moral hazard. For example, Sally may know that she will take some of her spare capacity offline during the premium period, creating an unexpected price rise in real-time. Ann protects herself from such events by requiring a larger price discount to accept a larger forward position.

Sally also expects tomorrow's premium price to be about $6/GB. However, as a natural seller, Sally is willing to sell at prices a few dollars below $6. Sally is happy to sell ahead an even larger portion of her expected production at higher prices. At prices well below $6, Sally is glad to buy ahead, knowing that the opportunity to sell tomorrow should reap profits. In the forward market, her offer must reflect the opportunity cost of selling the production tomorrow, which is about $6/GB. For Sally, like Ann, the forward market is about arbitrage and risk management. Her offers reflect opportunity cost, not marginal cost.

David anticipates that tomorrow's premium price will be about $6/GB. However, he is obligated to purchase his capacity needs tomorrow. He recognizes the possibility of demand shocks that could send the premium price to high levels. Thus, he adds a significant risk premium to his bids. He wants to buy a large share of his capacity needs unless the forward price is high. This preference is why David's net demand curve is significantly above Ann's and Sally's curves. His curve is similar in other respects: convex for negative quantities and concave for positive amounts.

All the curves are required to be piecewise-linear and decreasing. This language gives market participants enormous flexibility in expressing demand. The participant can approximate any continuous, decreasing demand. In this application, assuming that a participant's true preferences take this form is natural. An essential advantage of this form is that it implies unique prices and quantities, except in unlikely instances of no trade.

To find the clearing price, we add the individual demands in the quantity dimension, which yields the aggregate demand curve in Figure A2, focusing on the aggregate demand segment that includes the clearing price. The clearing price is where aggregate net demand is zero, a price of $6.11. The price is unique because the aggregate demand is continuous and strictly decreasing.

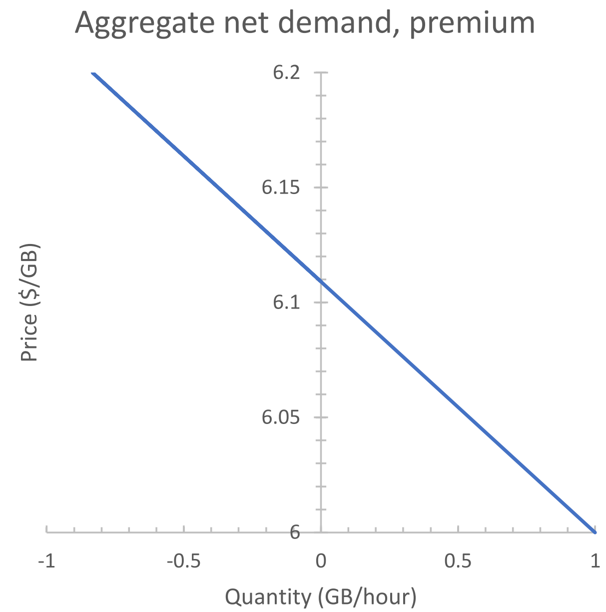

Fig. A2. Aggregate net demand curve around the clearing price for tomorrow's premium.

Finally, we determine the quantities by evaluating everyone's net demand at the clearing price, as shown in Figure A3. David buys at a rate of 3.8 gigabytes/hour; Sally sells at 3.6 gigabytes/hour; Ann sells at.3 gigabytes/hour. The net demand is zero, as required by market clearing.

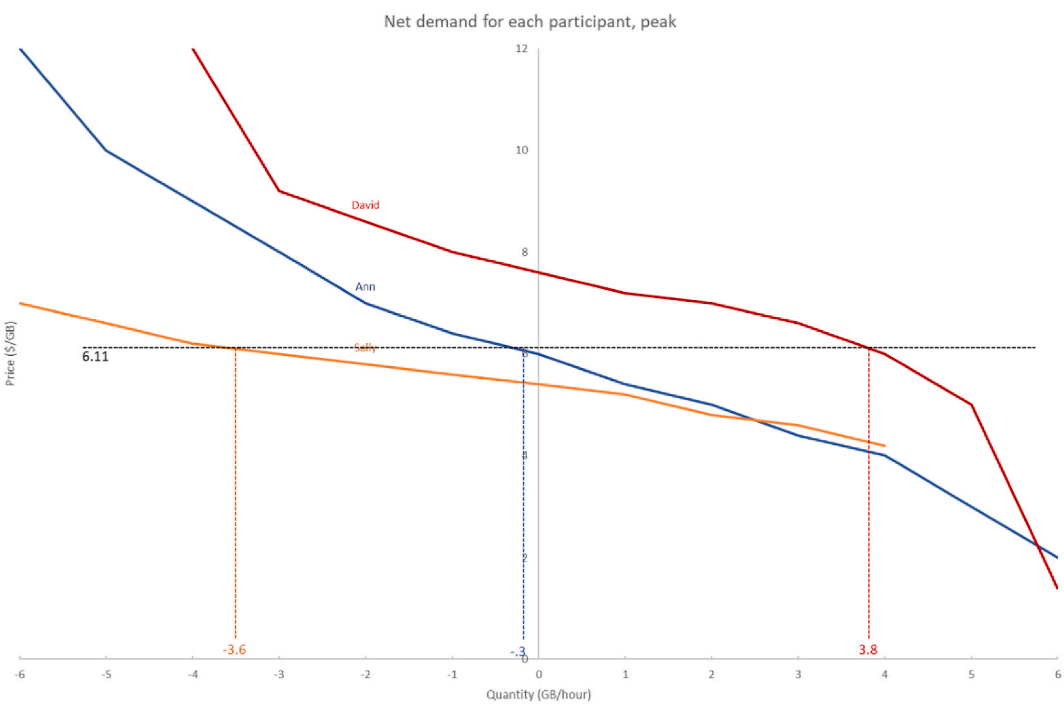

Fig. A3. Clearing quantities for each participant are uniquely determined from the clearing price.

Now, consider two products: regular and premium. Regular is for buyers who accept the slight possibility that their throughput will be rationed; premium is for buyers who find rationing unacceptable except in rare circumstances. How do the bid expression and market clearing generalize? With two products, our participants can bid on one or both products individually or on any linear combination of the two products.

Ann bids on the two products individually, and Sally and David bid on a linear combination of premium and regular, consistent with their objectives. Again, each order is a vector of quantity-price pairs, as shown inTable A2.

Table A2
Piecewise linear net demands (GB/hour) as a function of price ($/GB)

Quantity GB/hour	Price ($/GB)			
	Ann premium	Ann regular	Sally 50–50	David 60–40
−6	12	9	5.5	
−5	1	7	5.1	
−4	9	6	4.7	1.8
−3	8	5	4.5	8
−2	7	4	4.3	7.4
−1	6.4	3.4	4.1	6.8
0	6	3	3.9	6.4
1	5.4	2.4	3.7	6
2	5.	2	3.3	5.8
3	44	1.4	3.1	5.4
4	4	1	2.7	4.8
5	3			3.8
6	2	−1		0.2

Figure A4 shows the resulting demand curves. Each order is a piecewise linear decreasing curve. The participants can submit as many orders as they want, for individual products or any linear combination of products.

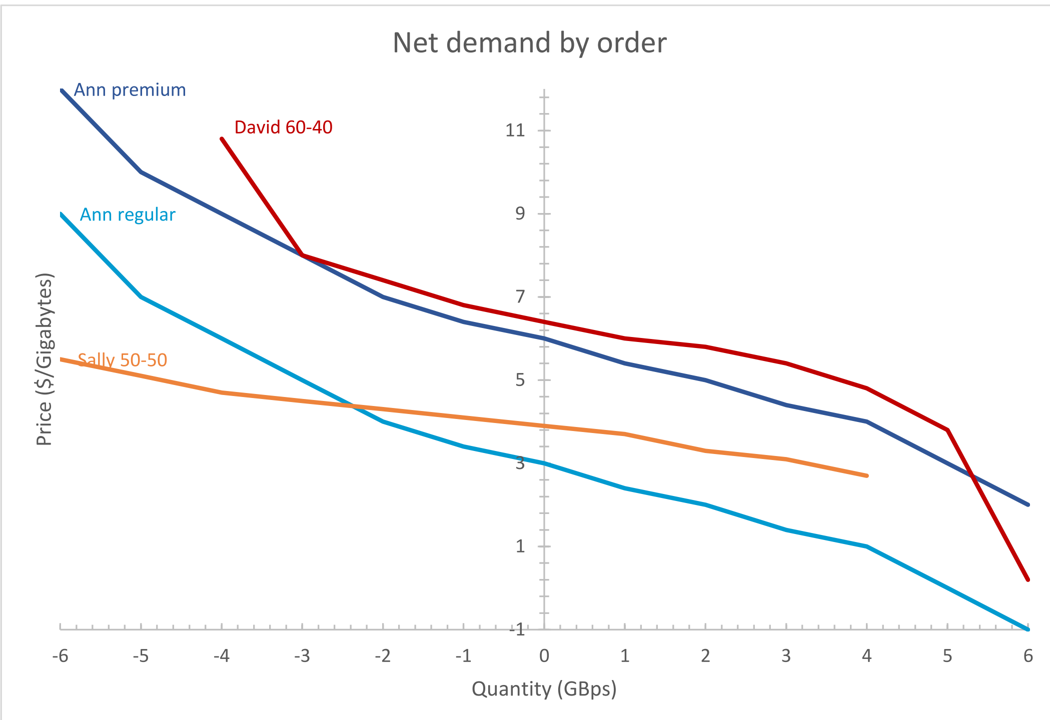

Fig. A4. Net demand by order.

Ann's regular order is her premium order shifted down by $3. She expects regular to clear at about $3. She wants to buy premium and regular communications at prices below $6 and $3 and sell them at higher prices.

David bids for a 60-40 split of premium and regular. His expected regular demand is 80 percent of his premium demand. Thus, the 60-40 split is consistent with his anticipated demand. He buys more premium communications because he appreciates its better throughput. He knows he will have to pay a risk premium for premium communications, but he is happy to do so to improve his communications.

Sally bids for a 50-50 split of premium and regular communications, slightly less than the 55-45 split she desires in real-time. She offers less premium communications because of the greater risk it entails. She knows shortages are possible, and the stringent premium quality is more challenging to satisfy. She expects a higher price on the premium communications she sells ahead.

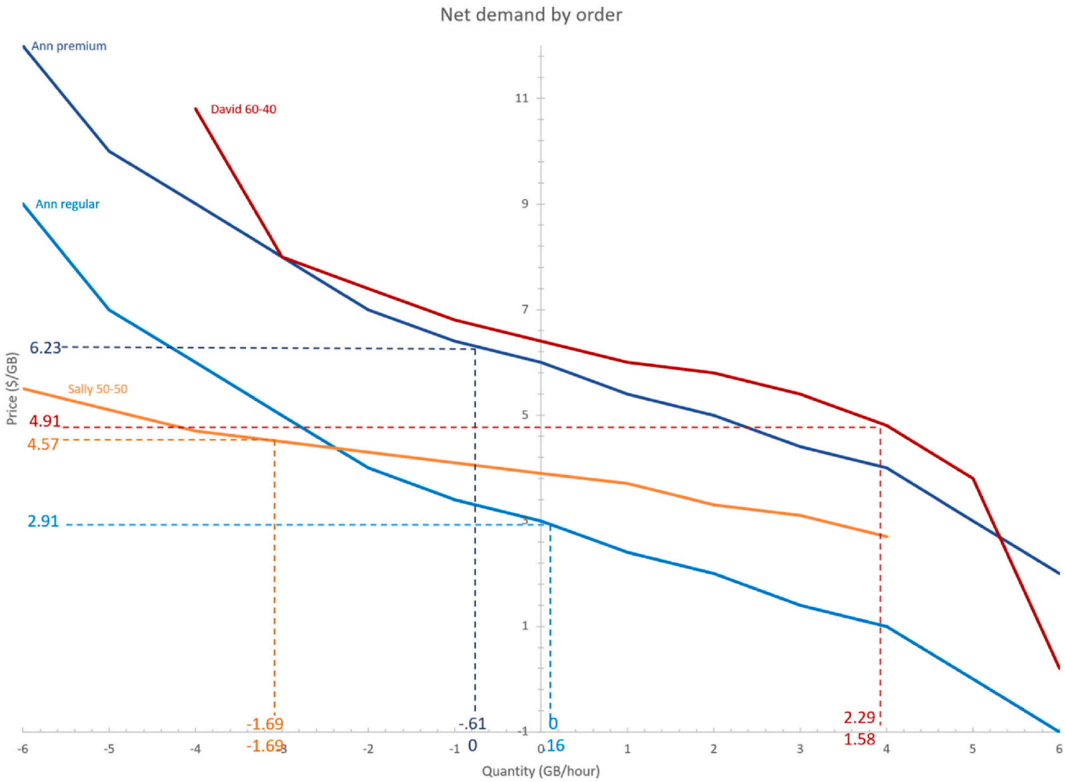

Fig. A5. Clearing quantities for each participant are uniquely determined from the clearing prices.

The products clear product-by-product. The number of prices is equal to the number of products. Market clearing involves finding two prices, premium and regular communications, that simultaneously balance supply and demand given the collection of orders. The clearing prices of $6.23 and $2.91 for premium and regular communications are displayed in Figure A5. These prices imply Ann sells 0.61 of premium communications and buys 0.16 regular. Sally sells 1.69 each of premium and regular communications. David buys 2.29 premium and 1.58 regular.

This example illustrates the beauty of the flow trading approach. The participants have enormous flexibility in expressing preferences. Then, given a collection of piecewise linear, decreasing demand curves, the market operator finds unique market clearing prices and quantities by solving a linear system. Larger problems with more products and more participants are solved similarly. The linear system to be solved is larger, but the computational needs are similar. Indeed, as discussed later, computation times tend to increase linearly with the number of products and orders.

References

Ausubel, L. M., Cramton, P., Pycia, M., Rostek, M., & Weretka, M. (2014). Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies, 81*(4), 1366–1400. https://doi.org/10.1093/restud/rdu023

Bertsekas, D. P. (2009). *Convex optimization theory*. Athena Scientific.

Black, F. (1971). Toward a fully automated stock exchange, part I. *Financial Analysts Journal, 27*(4), 28–35. https://doi.org/10.2469/faj.v27.n4.28

Bobbio, E., Brandkamp, S., Chan, S., Cramton, P. C., Malec, D., & Yu, L. (2023). *Resilient electricity requires consumer engagement*. University of Maryland Working Paper.

Borenstein, S. (2002). The trouble with electricity markets: Understanding California's restructuring disaster. *The Journal of Economic Perspectives, 16*(1), 191–211. https://doi.org/10.1257/0895330027175

Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning, 3*(1), 1–122. https://doi.org/10.1561/2200000016

Budish, E., Cramton, P., Kyle, A. S., Lee, J., & Malec, D. (2023). *Flow trading*. University of Maryland Working Paper.

Budish, E., Cramton, P., & Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics, 130*(4), 1547–1621.

California State Auditor. (2001). *California energy markets: Report 2001-0009*.

Cave, M., Genakos, C., & Valletti, T. (2019). The European framework for regulating telecommunications: A 25-year appraisal. *Review of Industrial Organization, 55*, 47–62.

Chao, H., Oren, S., & Wilson, R. (1986). Priority service: Pricing, investment and market organization. *The American Economic Review, 77*(5), 899–916.

Coase, R. H. (1972). Durability and monopoly. *The Journal of Law and Economics, 15*(1), 143–149.

Cramton, P. (2017). Electricity market design. *Oxford Review of Economic Policy, 33*(Article 4), 589–612.

Cramton, P. (2022). *Fostering resiliency with good market design: Lessons from Texas*. University of Maryland Working Paper.

Cramton, P., Brandkamp, S., Chao, H., Dark, J., Hoy, D., Kyle, A. S., Malec, D., Ockenfels, A., & Wilkens, C. (2024). *A forward energy market to improve resiliency*. University of Maryland Working Paper.
Cramton, P., & Doyle, L. (2017). Open access wireless markets. *Telecommunications Policy, 41*(5–6), 379–390. https://doi.org/10.1016/j.telpol.2017.01.003
Cramton, P., Gibbons, R., & Klemperer, P. (1987). Dissolving a partnership efficiently. *Econometrica: Journal of the Econometric Society*, 615–632.
Cramton, P., & Ockenfels, A. (2024). *Gradual and flexible procurement: A note on Germany's power station strategy*. University of Maryland Working Paper).
Jha, A., & Wolak, F. A. (2023). Can forward commodity markets improve spot market performance? Evidence from wholesale electricity. *American Economic Journal: Economic Policy, 15*(2), 292–330.
Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 1315–1335.
Myerson, R. B., & Satterthwaite, M. A. (1981). Efficient mechanisms for bilateral trading. *Journal of Economic Theory, 29*(2), 265–281.
Noam, E. (1998). Spectrum auctions: yesterday's heresy, today's orthodoxy, tomorrow's anachronism. Taking the next step to open spectrum access. *The Journal of Law and Economics, 41*, 765–790.
Pachler, N., Del Portillo, I., Crawley, E. F., & Cameron, B. G. (2021). An updated comparison of four low earth orbit satellite constellation systems to provide global broadband. In *2021 IEEE international conference on communications workshops (ICC workshops)* (pp. 1–7). IEEE. https://doi.org/10.1109/ICCWorkshops50388.2021.9473799.
Vayanos, D. (1999). Strategic trading and welfare in a dynamic market. *The Review of Economic Studies, 66*(2), 219–254.
Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance, 16*(1), 8–37.
Waddams, C. (2023). The UK energy price guarantee. *Intereconomics, 58*(1), 22–26.
Wolak, F. A. (2003). Measuring unilateral market power in wholesale electricity markets: The California market, 1998–2000. *The American Economic Review, 93*(2), 425–430.
Wolfram, C. D. (1999). Measuring duopoly power in the British electricity spot market. *The American Economic Review, 89*(4), 805–826.

Utilities Policy 18 (2010) 201–208

Contents lists available at ScienceDirect

Utilities Policy

journal homepage: www.elsevier.com/locate/jup

Virtual power plant auctions

Lawrence M. Ausubel*, Peter Cramton

Department of Economics, University of Maryland, College Park, MD 20742, USA

ARTICLE INFO

Article history:
Received 15 January 2010
Received in revised form
18 May 2010
Accepted 19 May 2010

Keywords:
Electricity auctions
Market design
Auctions
Clock auctions
Virtual power plants

ABSTRACT

Since their advent in 2001, virtual power plant (VPP) auctions have been implemented widely. In this paper, we describe the simultaneous ascending-clock auction format that has been used for virtually all VPP auctions to date, elaborating on other design choices that most VPP auctions have had in common as well as discussing a few aspects that have varied significantly among VPP auctions. We then evaluate the various objectives of regulators in requiring VPP auctions, concluding that the auctions have been effective devices for facilitating new entry into electricity markets and for developing wholesale power markets.

1. Introduction

Virtual power plant auctions are sales of electricity capacity which, rather than "physical" divestitures, are "virtual" divestitures by one or more dominant firms in a market. Instead of selling the physical power plant, the firm retains management and control of the plant, but offers contracts that are intended to replicate the output of the plant. Typically, these contracts are sold as divisible goods of varying durations, offered in periodic open and transparent auctions.

The motivation for and structure of a virtual power plant (VPP) auction is easiest seen by examining the Electricité de France (EDF) Generation Capacity Auctions, the world's first and longest-running series of VPP auctions. The EDF auctions began in 2001 as part of the regulatory *quid pro quo* for permitting EDF, the dominant electric utility in France, to proceed with the acquisition of a joint controlling stake in Energie Baden-Württemberg AG (EnBW), the fourth largest electric utility in Germany. The European Commission (EC) noted that EDF would be gaining joint control of one of the potential competitors particularly well placed to enter the French market, and the EC wished to require EDF to make available to other potential entrants a significant quantity of generating capacity in France. At the same time, given EDF's status as the largest nuclear producer in the world, the regulator recognized that physical divestment by EDF of its base-load nuclear plants would be undesirable in several respects. In particular, EDF had demonstrated a strong track record in the safety and security of its nuclear plants, and the public clearly benefited from economies of scale in EDF's management of nuclear plants. Consequently, the Undertaking agreed by the regulator and EDF in early 2001 provided for a virtual divestment by EDF of 6 GHz of French electricity capacity.

The VPP contracts offered in the EDF auctions are divided into two groups: base-load products and peak-load products. Each VPP product is an option contract for energy whose strike price approximates the variable cost of the respective energy. (For example, in the December 2009 auction, the strike prices of the base-load and peak-load VPP products were 10 €/MWh and 53 €/MWh, respectively.) As such, the base-load product is exercised essentially 24/7, whereas the peak-load product is exercised only a fraction of the time. Approximately 80% of the electricity capacity is offered as base-load products and approximately 20% is offered as peak-load products.[1] Within each of the two groups, a variety of durations would be offered: 3 months, 6 months, 12 months, 24

[1] In the early EDF Generation Capacity Auctions, there were actually three product groups: base-load VPP products; peak-load VPP products; and Power Purchase Agreement (PPA) products. The intention was for a total of 4000 MW to be offered of base-load VPP, 1000 MW to be offered of peak-load VPP, and 1000 MW to be offered of PPA. The PPA product was essentially a firm base-load product from November to March. Experience showed that the market had only limited demand for the PPA product. The parties eventually agreed to reconfigure the auctions so as to replace the 1000 MW of PPA product with 400 MW of VPP product. Thus, in the recent EDF auctions, the total quantity offered has been 4400 MW of base-load VPP product and 1000 MW of peak-load VPP product.

* Corresponding author.
E-mail address: ausubel@econ.umd.edu (L.M. Ausubel).
URL: http://www.ausubel.com, http://www.cramton.umd.edu

0957-1787/$ – see front matter
doi:10.1016/j.jup.2010.05.002

months and 36 months; all with the same starting date.[2] The principle followed for clearing is that all of the durations with a given starting date are treated equivalently for clearing: for example, 200 MW of a 3-month product offered in December, 200 MW of a 6-month product offered in December, and 200 MW of a 12-month product offered in December all count equivalently toward clearing. That is because each of these contracts has a start date of January 1st and, consequently, the sale of each of these contracts puts the same amount of electricity capacity in other parties' hands during the first quarter of the year. However, to the extent that the 3-month product is sold, an equivalent quantity will need to be auctioned again in March; to the extent that the 6-month product is sold, an equivalent quantity will not need to be auctioned until June; and to the extent that the 12-month product is sold, an equivalent quantity will not need to be auctioned until the following December.

The first EDF auction was conducted in September 2001; and, as of this writing, there have been 34 quarterly auctions successfully held. Meanwhile, the VPP auction has proven popular with regulators throughout Europe. The basic mechanism has been replicated for: Electrabel in Belgium; Nuon in the Netherlands; Elsam in Denmark; Endesa and Iberdrola, in combined auctions, in Spain; REN and EDP, in combined auctions, in Portugal; and E.ON and RWE, in separate voluntary auctions, in Germany. A similar structure was used in the Texas Capacity Auctions, in the US; and was planned in connection with the Exelon-PGE merger, in the US.[3] In addition, the so-called "gas release programme auction" — the natural gas counterpart of the VPP auction — has been utilized in Germany, Austria, France, Hungary and Denmark.

In requiring VPP auctions, regulators may be attempting to further any or all of the following objectives:

- Facilitating entry into the electricity market by assuring the availability to new entrants of electricity supplies on the high-power grid;
- Promoting the development of and adding liquidity to the wholesale electricity market; and
- Reducing market power in the spot electricity market.

This paper proceeds as follows. Section 2 describes and explores the various design choices that virtually all VPP auctions have had in common, while Section 3 focuses on a few aspects of the auction design that have varied significantly among the VPP auctions to date. Section 4 considers whether VPP auctions have been an effective tool for promoting each of the pro-competitive objectives listed above, and Section 5 concludes.

2. Commonalities in the design choices for VPP auctions

Virtually all of the VPP auctions that have been adopted in practice have made the same general design choice: a simultaneous ascending-clock auction with a discrete round structure. In this section, we describe the simultaneous ascending-clock auction and we explore the reasons for the unanimous choice.[4]

2.1. Simultaneous ascending-clock auction with discrete rounds

In the ascending-clock auction with discrete rounds, the following basic procedure is typically used:

- The auctioneer pre-announces an available supply, S, in the auction, which may be subject to a reserve price or an increasing supply curve;
- The auctioneer announces to bidders an interval of prices, $[\underline{p}_t, \overline{p}_t]$, effective for round t;
- Each bidder i simultaneously and independently submits its demands $q_i(p)$ for prices $p \in [\underline{p}_t, \overline{p}_t]$, during round t, where $q_i(p)$ is constrained to be a downward-sloping demand curve;
- Following round t, the auctioneer calculates the aggregate demand $AD \equiv \sum_{i \in I} q_i(\overline{p}_t)$;
- If $AD > S$, then the aggregate demand AD is disclosed to the bidders and the auction progresses to round $t+1$, in which an interval of prices $[\underline{p}_{t+1}, \overline{p}_{t+1}]$, where $\overline{p}_{t+1} > \underline{p}_{t+1} = \overline{p}_t$, is effective; and
- If $AD \leq S$, then the auction concludes at a clearing price of $p^* \in [\underline{p}_t, \overline{p}_t]$, where p^* is typically selected to be the smallest p such that $\sum_{i \in I} q_i(p) \leq S$.

When the ascending-clock auction involves multiple products, they are typically auctioned simultaneously. Products may be in the same product group or in distinct product groups. When products are in the same product group, it is possible for bidders to "switch" from one product to another as prices ascend; while when products are in distinct product groups, they are auctioned independently (but simultaneously). For example, in many of the auctions, base-load products of different durations have been assigned to the same product group, while peak-load and base-load products have been assigned to different product groups. The rationale for this grouping has been that base-load products of different durations are generally viewed as substitutes, while base-load and peak-load products are generally viewed as complements. As such, a bidder may wish to shift its demand among the different base-load products as prices evolve, but probably will not need to shift its demand between base-load and peak-load products.

2.2. Dynamic vs. sealed-bid

By contrast, in the standard sealed-bid auction, bidders have a single opportunity to submit demand curves $q_i(p)$ that cover the entire possible range of prices.[5] They do not receive any feedback about the bids of other bidders until the auction has concluded. Based on the single round of sealed-bid submissions, the auctioneer determines the clearing price p^* to be the smallest p such that $\sum_{i \in I} q_i(p) \leq S$ (or the largest p such that $\sum_{i \in I} q_i(p) \geq S$). Each bidder i wins the quantity $q_i(p^*)$ and pays either p^* per unit (uniform-price auction) or the amount of its winning bid (pay-as-bid auction), depending on the exact auction format.

[2] Subsequently, a 48-month product has been added to the base-load VPP product group. Moreover, the September auction has offered additional products: in addition to the usual array of products with start dates of 1 October (one month after the auction), there are 2-month, 12-month, 24-month, 36-month and 48-month (base-load only) products with start dates of 1 November (i.e. two months after the auction).

[3] However, the Exelon-PGE merger failed to receive the approval of New Jersey regulators, and merger efforts were ultimately abandoned.

[4] Additional details and discussions relating to dynamic clock auctions may be found in Ausubel and Cramton (2002, 2004), Ausubel et al. (2002), and Ausubel (2004). An exploration of the relationship between VPP prices and spot prices in the French market can be found in Armstrong et al. (2007).

[5] Often, in sealed-bid auctions, bidders are permitted to submit multiple bids, each for a given quantity of electricity and at a given price. The reader should observe that, when bids of a bidder are expressed in the latter form, they may be combined together to form an inverse demand curve, and the expression of an inverse demand curve is almost equivalent to the expression of a demand curve. Thus, the latter form is almost equivalent to the submission by bidders of demand curves.

For virtual power plants, dynamic auction formats offer at least four decisive advantages over sealed-bid auction formats. First, dynamic auctions offer the greatest transparency and, by contrast, sealed-bid auctions are comparatively opaque. Recall that VPP auctions are frequently invoked as competition remedies for facilitating entry into markets with dominant firms; consequently, it is important to the credibility and success of the programs for competitors, regulators and the public to be able to see that the auctions are conducted fairly and in accordance with the published rules. The transparency of dynamic auction formats is thus an important property favoring their adoption for VPP auctions. Second, an ascending-clock auction is a particularly simple and effective format for obtaining price discovery. Since another frequent objective of VPP auctions is to jumpstart the development of wholesale power markets, the promotion of price discovery (which, in turn, facilitates wholesale power transactions outside the auction) is another valuable feature of ascending-clock auctions. Third, in trying to explain why dynamic auctions are growing in popularity relative to sealed-bid auctions, the literature has observed that bidders will be reluctant to reveal their valuations truthfully in an auction where the seller may have the opportunity subsequently to use the information against the bidders. By contrast, a dynamic auction avoids this problem, as it does not require the high-value bidders to reveal their true valuations — the bidding stops as soon as the aggregate demand becomes equal to supply. Again, this issue is likely to be important in VPP auctions, as the seller is a dominant firm, and the bidders are potential entrants. Fourth, the ascending-clock auction format scales particularly well to a simultaneous auction of multiple products, which are frequently present in VPP auctions. By contrast, independent sealed-bid auctions perform particularly poorly when substitutes — for example, base-load products of different durations — or complements — for example, base-load and peak-load products — are auctioned together.

Given these advantages, it is not at all surprising that essentially all virtual power plant auctions to date have utilized some variation on a simultaneous ascending-clock auction.

2.3. *Discrete rounds vs. continuous bidding*

Although in theory one can imagine implementing an ascending-clock auction in continuous time, this is hardly ever done in practice in auctions of high-valued items. VPP auctions inevitably use discrete rounds for at least three important reasons. First, communication is rarely so reliable that bidders would be willing to be exposed to a continuous clock. A bidder would find it unsatisfactory if the price clock swept past the bidder's willingness to pay because of a brief communication glitch. Discrete rounds are robust to communication problems. Discrete rounds have a bidding window of significant duration, rarely less than ten mintues and often a half-hour or longer. This window gives bidders time to correct any communication problems, to resort to back-up systems, or to contact the auctioneer and have the round extended. Second, bids need to be legally-binding commitments in order for an auction process to work as intended. This implies that bidders need to be given sufficient time to reflect upon, carefully enter, check and submit their bids, if bidders are going to be held to their bids. Third, a discrete-round auction also improves price discovery by giving the bidders an opportunity to reflect between rounds. Bidders need time to incorporate information from prior rounds into a revised bidding strategy. This updating is precisely one of the sources of price discovery and its associated benefits.

It is only in sequential *descending* clock auctions (Dutch auctions) that a nearly continuous bidding process is used. This is seen in Dutch flower auctions, many fish auctions, and US tobacco auctions since 2003. All of these auctions are conducted on-site (avoiding communication difficulties) and they all involve descending clocks (reducing the role for price discovery within the auction).

2.4. *Divisibility of the product*

Given that electricity is nearly a perfectly-divisible good, it is natural for the auction process to treat it as highly divisible. Thus, many VPP auctions (e.g. France, Belgium and Denmark) have used minimum bidding units of 1 MW, in auctions where anywhere from 100 to more than 1000 MW of contracts are offered. The initial Spanish VPP auctions used bidding units of 2 MW — that was because the auctions were conducted jointly for Endesa and Iberdrola, and so the minimum bid was 1 MW attributable to each seller. Later, the bidding unit was raised to 10 MW (5 MW for each seller). In the E.ON VPP auction, the minimum positive bid was 5 MW, but above that, the bidding unit was 1 MW; the minimum was only to establish a minimum scale where it would be worth setting up contractual arrangements with a winner.

2.5. *Activity rule*

To promote price discovery, activity rules are generally imposed in ascending-clock auctions. In an ascending-clock auction for a single product, the prevalent activity rule takes the simple form of a monotonicity constraint: each bidder's quantity demanded is not permitted to increase as the price increases, consistent with downward-sloping demand curves. Without the monotonicity constraint, a bidder might hide as a "snake in the grass" — grossly understating demands at low prices and then jumping in with large demands near the end of the auction. Widespread use of a snake-in-the-grass strategy would undermine the very purpose of utilizing a dynamic auction.[6] A monotonicity constraint prevents this form of strategic behavior, thus encouraging better price discovery and facilitating rapid convergence to equilibrium.

In situations with multiple goods that have relatively independent demands or are complements, a monotonicity constraint is often applied independently to each good. However, in situations where two or more products are close substitutes, applying monotonicity constraints independently to each good may be overly restrictive; it is natural for the bidder to want to switch to the product with the more attractive price. This would be excluded by the simplest application of independent monotonicity constraints.

A common approach is to organize different durations of the same type of contract into product groups. Since the goods within a group are denominated in comparable units (MW of power), the activity rule applied to all products within a group can simply be a monotonicity constraint on the sum of the demands for the respective products. This approach was utilized in the French and Belgian VPP auctions; it permits bidders to substitute among 3-month, 6-month, and 12-month contracts, etc., on a one-to-one basis. A variation on this approach has been utilized in the Spanish auctions: there, contracts of different durations are compared according to the total number of months, so that there, bidders can

[6] One motivation for a bidder to use a "snake-in-the-grass" strategy is to avoid conveying information to rivals in an environment where bidders exhibit interdependent values. If each bidder's estimate of value is based in part on rivals' information, one bidder demanding large quantities might induce her rivals to raise their value estimates and bid more aggressively. A second motivation for a bidder to use a snake-in-the-grass strategy arises from budget constraints. The bidder holds back on bidding for the good she wants most dearly, instead bidding for the goods her rivals want, in the hopes of exhausting the competitors' limited budgets. The bidder then shifts to bidding on her true interests late in the auction, now facing weakened competition for these goods.

substitute among 3-month, 6-month and 12-month contracts, on a 4-to-2-to-1 basis.

2.6. Information disclosure during the auction

In an ascending-clock auction, there are many possible policies for information disclosure during the auction. With respect to the level of aggregation or disaggregation of bidders' demands, one could disclose the aggregate demand for each product, disclose each individual bidder's demands anonymously, or disclose each individual bidder's demands identified by bidder. With respect to the information provided about demand in the price interval $[\underline{p}_t, \overline{p}_t]$, one could disclose demand at the end-of-round price $\overline{p}_t$ only, or one could disclose demand at all prices in the interval $[\underline{p}_t, \overline{p}_t]$.[7]

Reporting only the aggregate demand for each product at the end-of-round price, after each round, has been viewed as striking a comfortable balance between information useful for price discovery and information that facilitates collusion. In VPP auctions, the aggregate demand for a product contains most of the information needed for price discovery. If, instead, the auctioneer revealed the individual demands of each bidder, this detailed information could be used to coordinate reductions in demands at low prices. For example, the bidders might cooperatively reciprocate the quantity reductions of competitors, and attempt to punish those who do not reciprocate by shifting quantity toward products most desired by the non-reciprocating bidder. Consequently, in essentially all VPP auctions, the determination has been made to report only the aggregate demands for the products after each round.

2.7. Internet auction

Best practice for conducting auctions of high-valued items, today, is by internet-based software. Gathering the bidders together in a single location would both be unnecessarily disruptive to participants, whose offices are located across wide geographic regions, and be unnecessarily conducive to collusion. As such, essentially all VPP auctions have been conducted online on the Internet.

2.8. Frequency of the auction

Most VPP auctions have been conducted at frequent intervals. The EDF, Electrabel, Elsam and RWE auctions were all scheduled as quarterly auctions; while the Texas Capacity Auctions were conducted about five times per year. The Endesa–Iberdrola VPP auctions were also initially held quarterly, but they were later changed to be semi-annual auctions.

As devices for facilitating competition in the market of a dominant firm, frequent VPP auctions are helpful in offering entrants frequent opportunities to bid for assured supplies within the market. Entrants can buy electricity capacity when they need it, and they can adjust their purchases according to the penetration they achieve in the market. As devices for adding liquidity to the forward market, frequent releases of supply are also useful.

Sellers also often tend to value holding frequent auctions. By contrast, offering a significant fraction of a firm's capacity on a single date subjects the firm to a significant amount of market risk; sellers tend to prefer spreading out the sales over several auction dates so as to reduce the risk associated with market fluctuations. Moreover, there tends to be greater liquidity (and greater demand by bidders) for products of relatively short duration (3-month to 24-month contracts), as compared to longer-term contracts. There also tends to be limited appetite for purchases of products on a given date, which can easily be exhausted by offering a large supply of contracts on a one-off basis on a single date. Thus, sellers find frequent auctions much more palatable, which helps to explain why negotiated settlements with regulators often tend to include relatively frequent auctions.

However, it should be observed that the implementation of VPP auctions at frequent intervals sets apart the "virtual" divestiture from a "physical" divestiture, which would typically be the one-off sale of the entire useful life of a generating asset on a single date. This substantive difference between virtual and physical divestitures will be explored further in Section 4.

3. Differences in the design choices for VPP auctions

While essentially all virtual power plant auctions to date have followed a common basic structure, described in Section 2, there have also been significant differences in the design choices made. This section considers some of the differences.

3.1. Fixed supplies of one or more duration versus supply flexibility

The VPP auctions to date have taken three divergent approaches to the durations of VPP contracts. In some (as exemplified by the French VPP auctions), several different durations with the same starting date are offered in each auction, with the clearing condition based only on the total quantity sold and no preconditions on the quantities sold of any particular duration. In others (as exemplified by the Danish VPP auctions), a limited set of durations is offered in each auction and only a fixed predetermined quantity of each is sold. And in others, only a single duration is offered in the auction.

In the case of the French auctions (as well as the auctions in Belgium and Spain), it was recognized that different bidders might prefer buying different durations. The view taken was that the regulators' interest was only in the aggregate flow quantity of VPP contracts in the hands of parties other than the dominant firm at any moment in time, and not in the duration that these contracts would take. Meanwhile, neither EDF nor the regulators had a reliable method for predicting the demands for the various durations — other than through the auction itself — and the relative demands for the various durations might change from auction to auction, depending on which bidders choose to participate and their respective needs. By way of contrast, there existed good methodology for developing the "term structure" of relative valuations for the contracts of various durations that would make the seller indifferent between selling one duration or another.

Observe that if both the quantities and the relative prices of the various durations were allowed to be determined endogenously, then the entire system would be underdetermined. For example, suppose that it was decided that 500 MW of base-load power would be sold as 3-month or 12-month contracts, and that no quantity relationship or price relationship would be imposed on sales of the two contracts. Then observe that one possible outcome would be prices p_3 and p_{12} such that 500 MW of the 3-month contract and 0 MW of the 12-month contract were demanded by bidders. A second possible outcome would be prices p'_3 and p'_{12} such that 250 MW of the 3-month contract and 250 MW of the 12-month contract were demanded. And a third possible outcome would be prices p''_3 and p''_{12} such that 0 MW of the 3-month contract and 500 MW of the 12-month contract were demanded. Then, under ordinary demand conditions for substitutes, we would expect that $p_3 < p'_3 < p''_3$ and $p''_{12} < p'_{12} < p_{12}$. That is, the auction

[7] One could also elect not to disclose any demand information after each round, other than the fact that aggregate demand exceeds supply and so the auction remains open. But this would run opposite to the motivation for using an open dynamic auction, and so this policy of nondisclosure is seldom taken.

Table 1
Quantity sold and final prices in June 2009 EDF VPP auction.

Duration of base-load product	3-Month	6-Month	12-Month	24-Month	36-Month	48-Month	Aggregate total
Quantity sold	95 MW	40 MW	70 MW	125 MW	50 MW	100 MW	480 MW
Final price per month	€19,500	€25,556	€28,811	€31,605	€33,217	€34,706	

outcome would not be pinned down at all unless the quantities to be sold of the different durations were pre-specified or if the price relationship was pre-specified.

Since the composition of different durations was intended to be market driven and since the term structure of prices was reasonably well understood, the decision was made that the prices of the various product durations within a group would be linked together and would increase in lockstep. (However, the prices associated with different product groups — base-load versus peak-load — move independently of one another.) Before the start of the auction, and under the supervision of a trustee, the seller determines an "indifference table" expressing the price differentials (i.e. a yield curve) amongst the various products within a group that would make the seller indifferent between selling one product or another. With two product groups containing six and five products, respectively, there are effectively just two degrees of freedom (and two price "clocks"), although eleven prices in total. The clearing condition is then that the aggregate demand for each product group is to be no greater than the total supply offered. The auction itself then determines endogenously the distribution of sales across the various durations.

Table 1 illustrates the success of this approach by providing the results, with regard to both quantities and prices of the base-load products, in the June 2009 EDF auction. The last row of Table 1 shows the indifference table that was used in the auction. Prices prior to the final round were additive transformations of this curve: for example, the end-of-round prices of Round 1 were (€17,300 €23,356 €26,611 €29,405 €31,017 €32,506), respectively, for the six different durations. In general, there may be some minor concerns that the seller might attempt to manipulate the indifference table to its advantage: for example, if the seller believed that bidders favoring 24-month or 36-month contracts were more effective competitors than bidders favoring shorter-term contracts, then the seller might price the longer contracts disadvantageously. However, apart from the obvious difficulties for the seller in obtaining sufficient information to use such a strategy, making such manipulation implausible, observe that the results displayed in Table 1 are strongly suggestive of a fair indifference table. Aggregate demand for each of the six durations was no lower than 8% and no greater than 26% of the total demand, aggregated over all durations.

The June 2009 EDF auction also included 556 MW of "advance sales" of products to be offered in the September auction: ten products of various durations with starting dates of 1 October or 1 November 2009 (not shown in Table 1). These were similarly offered with a yield curve of indifference prices; and a positive quantity of each of these ten products was sold.

In some other series of VPP auctions, multiple durations are offered to bidders, but only in fixed supplies. Table 2 illustrates the approach that has been taken in the Danish VPP Auctions. Contracts of 3-month, 12-month and 36-month durations are offered according to a planned schedule, in fixed supplies of 100 MW or 200 MW in a given auction.

Meanwhile, in some other VPP auctions (e.g. the Netherlands, RWE Germany and Portugal auctions), only a single product duration was generally offered to bidders.

The approach of supply flexibility appears the most desirable, for three reasons. First, in terms of the objectives of facilitating the obtaining of supply by new entrants and of increasing the liquidity of wholesale markets, the extra flexibility is highly desirable. New entrants are better able to obtain quantities of electricity capacity over time that match their needs; while new liquidity will gravitate to durations that are in the greatest need of liquidity in the wholesale market. Second, value is maximized among sellers and bidders by offering flexibility in duration: if there are greater gains from trade at a particular duration, the auction will shift sales toward that duration. Third, the probability of a product failing to sell (due to receiving bids less than the supply) is minimized, improving the likelihood that the regulatory objectives of the VPP auction program are met.

By the same token, the approach of offering multiple durations, each in fixed quantities, appears to be superior to offering only a single duration. Given the heterogeneity of bidders, it is unlikely that a "one size fits all" contract would meet entrants' needs or maximize gains among sellers and bidders. Additional durations and additional flexibility will generally be beneficial.

3.2. Structure of bid submissions

The VPP auctions to date have also taken three divergent approaches to the exact structure of bid submissions. In many (for example, the French, Belgian, Spanish and E.ON German VPP auctions), bidders are permitted in round t to submit essentially arbitrary non-increasing step functions of quantities associated with the interval of prices, $[\underline{p}_t, \overline{p}_t]$. In some auctions, bidders are permitted in round t to submit step functions of quantities with a single reduction in the interval $[\underline{p}_t, \overline{p}_t]$ ("exit bids"). And in a few auctions (for example, the Dutch and Danish VPP auctions), bidders are permitted to submit quantities at only the single price $\overline{p}_t$, and the resulting "overshoot" is resolved by having bidders re-bid in a final sealed-bid round.

The approach of having bidders re-bid in a final sealed-bid round generally achieves poor results relative to the objectives of efficiency or revenue maximization. The reasoning is as follows: suppose a situation where the true clearing price is $½(\underline{p}_t + \overline{p}_t)$, the midpoint of the interval of prices effective in round t. Then the auction will attract insufficient demand in round t at the price $\overline{p}_t$, and the bids of round t will be re-bid in a final sealed-bid round. Bidders, learning the "bad news" that there was insufficient demand at $\overline{p}_t$, will (regardless of their preexisting assessments of value) tend to bid close to the minimum allowable amount of $\underline{p}_t$ in the re-bidding. Thus, the contracts will tend to be allocated randomly among the remaining bidders rather than allocated efficiently to the bidders with the highest valuations, and the revenue per contract will tend to be approximately $\underline{p}_t$ — unambiguously less than $½(\underline{p}_t + \overline{p}_t)$.

By the same token, allowing bidders to submit essentially arbitrary non-increasing step functions of quantities associated with the current interval of prices will tend to produce the true clearing price, with bids dispersed according to the bidders' underlying valuations rather than clustered at the minimum possible price. Thus, the design used in France, Belgium, Spain and for E.ON in Germany tends to produce more efficient and higher revenue outcomes. Meanwhile, the approach of allowing bidders to submit step functions of quantities with a single reduction is a partway measure that also improves upon the approach of re-bidding

Table 2
Quantities offered in Danish VPP auctions, May 2008–Nov 2009.

Duration of product offered	3-Month	12-Month	36-Month
Quantity offered in Aug 2008 auction	200 MW		100 MW
Quantity offered in Nov 2008 auction	100 MW	200 MW	
Quantity offered in Feb 2009 auction	100 MW		
Quantity offered in May 2009 auction	100 MW		
Quantity offered in Aug 2009 auction	100 MW		100 MW
Quantity offered in Nov 2009 auction	100 MW	200 MW	

in the final round but sacrifices some of the efficiency and revenues achievable using arbitrary step functions.

3.3. *Reserve prices*

The VPP auctions conducted to date have varied in their reserve-price policies. The French VPP auctions have not utilized any reserve price for the base-load and peak-load products, but this was accompanied by a confidence that aggregate demand in the auction would far outstrip the supply.[8] Indeed, in the typical EDF auction, there has been approximately a four-to-one ratio between the aggregate demand in Round One and the supply. Most of the other VPP auctions have utilized some form of announced or secret reserve price, but this was accompanied by the recognition that demand in many of the other markets was much weaker and that the seller needed the protection of a reserve price in the event of insufficient demand.

An announced reserve price can be implemented very simply in an ascending-clock auction, by starting the price clock at the reserve price. A secret reserve price is typically implemented under the supervision of a trustee or monitor, who assures that the reserve price has been fixed before the bidding starts. A given product does not clear until aggregate demand is less than or equal to the supply *and* the reserve price is reached. If the aggregate demand is less than or equal to the supply but the auction remains open, bidders can infer that the reserve price is the level at which the auction ultimately closes; but otherwise the "secret reserve" price stays undisclosed.

Observe that a reserve price is a useful instrument for addressing limited competition within the auction. It does this in two ways. First, it reduces the incentive for collusion by limiting the maximum gain from collusion. Bidders must pay at least the reserve price no matter how effective their collusion. Second, an appropriately chosen reserve price guarantees that the seller receives a significant fraction of value, even when competition is weak.

A generalization of a reserve price is for the auction to utilize an increasing supply curve. In a clock auction, a supply adjustment is most easily accomplished by specifying an explicit upward-sloping supply curve. This has the effect of expanding the quantity offered for sale when there is ample competition, but reducing the quantity offered (and implicitly introducing a reserve-like mechanism) when there is insufficient competition within the auction.

3.4. *Information disclosure after the auction*

The VPP auctions conducted to date have varied substantially in their post-auction information-disclosure policies. In many (for example, the French, Belgian and Spanish VPP auctions), the same information that becomes available to winning bidders during the auction is also made available to the general public shortly after the auction. The disclosed information includes the prices and aggregate demands for each product after every round, including the final round. In some other VPP auctions, the only information that is announced publicly is the final price and quantity.

A policy of widespread disclosure is preferable for several reasons. First, it facilitates participation by new entrants, by putting them on a level playing field with past participants. Second, it helps to assure a high level of transparency in the auction process. Finally, the disclosure enhances the secondary market.

4. VPP auctions as tools for promoting competition and liberalization

The most common motivation for virtual power plant auctions has been to promote competition in and the liberalization of electricity markets with one or more dominant firms. In this section, we explore and evaluate the possible pro-competitive effects of VPP auctions. Commentators have suggested at least three mechanisms by which VPP auctions may promote competition and liberalization:

- They may facilitate entry into the electricity market by assuring the availability to new entrants of electricity supplies on the high-power grid;
- They may promote the development of and add liquidity to the wholesale electricity market; and
- They may reduce market power in the spot electricity market.

The first two mechanisms have been foremost in the minds of regulators and are probably the most important. For example, in the merger procedure leading up to the EDF auctions, the European Commission (2001) wrote: "Access to generation capacity in France would only realistically be possible if EDF granted such access since EDF is the main generator in France." (paragraph 34). In assessing the competitive situation in 2001, the EC concluded: "Newcomers have only marginal chances to purchase electricity in the framework of trading in France" (1.3.2.2); "Newcomers face difficulties when entering the French market via imports" (1.3.2.3); and "The overwhelming position in electricity generation in France allows EDF to outbid competitors trying to enter the French market" (1.3.2.4).

In requiring VPP auctions as a *quid pro quo* for allowing EDF to take a joint controlling interest in EnBW, the EC believed that the new auctions would facilitate new entry and competition: "The access to generation capacity will enable foreign suppliers to become active on the market for supply to eligible customers to a significant extent." (paragraph 107). "Furthermore, German suppliers will also be able to gain a foothold in France and thus become sufficiently strong in France in order to cope with EDF's potential for retaliation resulting from its presence in Germany." (paragraph 108). "Finally, the access to generation capacity in France will put foreign suppliers in a better position regarding Pan-European supply contract since they will be able to supply customers with eligible production sites in France through a VPP contract with EDF." (paragraph 109).

The various VPP auctions appear to have been generally successful, operating primarily through the first two mechanisms. One important data point is the development of the wholesale electricity market in France. In 2001, any wholesale electricity market was close to nonexistent in France — to the point that, for the setting of reference prices in the early EDF auctions, the price data was taken from the German wholesale market (the French data being too thin and lacking in meaning). However, after eight years of VPP auctions, the French market is now generally considered to be the third most active electricity wholesale market in Europe.

[8] A reserve price was introduced in the French auctions of 2003–2005 for the Power Purchase Agreement (PPA) product, after it became clear that the demand for the PPA product was much lower than for the base-load and peak-load VPP products. When the PPA product was discontinued in 2006, the reserve price was also discontinued.

Various European utilities today view participation in VPP auctions as an important element of their pan-European strategies. For example, Iberdrola (itself required to sell in the Spanish VPP auctions) recently trumpeted the fact that it had successfully acquired 1500 MW of capacity in 2008 in VPP auctions in Germany, France and Portugal (Iberdrola, 2009). While European utilities have been consolidating, their operations outside their principal markets have been expanding, partly due to the access to capacity afforded by VPP auctions.

The third mechanism has been emphasized, for example, by Christian Schultz (2005). Physical or virtual divestitures by dominant firms have the potential to reduce market power in the spot market by creating less concentrated market shares in generating capacity. Schultz argues that VPP auctions, as typically implemented, diminish spot market power much less than is possible, since the contracts are relatively short-lived (as compared to physical divestiture) and the auctions are generally frequent. Schultz would therefore prefer unstaggered VPP contracts of long duration (or physical divestitures).

Our assessment is that VPP auctions as currently practiced are not oriented toward making major reductions of concentration levels in spot markets. VPP contracts are intended to be relatively long-term contracts, and they are intended to be bought by competitors of the dominant firm(s). This means that the demand for VPP contracts is relatively limited. Consequently, VPP auctions as currently practiced must involve a relatively small fraction of electricity capacity in the given market. For example, in France in 2001, EDF accounted for greater than 80% of the overall electricity market, while the VPP auctions have never sold more than 10% of total generating capacity. In Spain, Endesa and Iberdrola together accounted for greater than 60% of the overall electricity market, while the VPP obligations were less than 6% for Endesa and less than 5% for Iberdrola (Federico et al., 2008).

Thus, the current magnitudes of electricity assigned to the VPP auctions are insufficient to have a major impact on concentration levels in spot markets. Moreover, there may be no practical way to increase the capacities subject to VPP auctions; even at current levels of sales, many of the VPP auctions (outside France) have bumped against the reserve prices.

While VPP auctions are effective devices to enable entrants to gain footholds in markets with dominant firms and for developing wholesale markets, they are thus ill-suited for making major changes in spot market concentrations. By contrast, forward markets are effective devices for reducing market power in the spot electricity market (Ausubel and Cramton, in this issue). More than anything, what distinguishes the forward auctions useful for correcting the spot market from VPP auctions is that the buyers of contracts in such forward markets are principally the load (hedging the spot market), while the buyers of VPP contracts are ideally competitors (enabling new entry). Note also that, while a VPP auction obligation is normally placed on a dominant firm, the forward trading by suppliers in the envisioned forward market should extend to all generators.

In the longer term, one could easily imagine the current VPP auctions enlarging and evolving in the direction of larger auctions that take on the dual role of facilitating entry by new suppliers and yielding forward sales from suppliers to load. But, for this to occur, the regulatory structure will need to evolve in a direction where suppliers and load are both given incentives or obligations to engage in forward trading.

5. Conclusion

We have reviewed the structure of virtual power plant auctions that began in 2001 and have subsequently spread widely in use. We have seen the aspects of the auction design that are common to essentially all VPP auctions, and we have seen the aspects that differ among the various auctions. We have also seen that VPP auctions are effective devices for facilitating new entry into electricity markets and for developing wholesale markets, while they have not been oriented toward making substantial reductions of concentrations in spot markets.

One important reason for evaluating the various mechanisms by which VPP auctions can promote competition is that it provides insights into the appropriate duration of VPP contracts and the appropriate frequency of auctions. If the primary objective was to equalize market shares in the spot market, then VPPs might be designed to replicate physical divestitures as closely as possible: contracts would be extremely long term, and they would be sold in one-off auctions. However, such timing would be antithetical to the principal objectives of facilitating new entry and of developing wholesale markets. Entry would be facilitated only at the time of the one-off auction; and later arrivals might find themselves lacking any mechanism for obtaining capacity. Meanwhile, liquidity would not be added for those contracts (such as 3-month and 12-month contracts) that can most plausibly become actively-traded products in wholesale electricity markets. As such, one can expect that most VPP contracts offered will continue to be in 3-month to 36-month durations, and that most VPP auctions will continue to occur quarterly.

Given their success, VPP auctions will deservedly continue to receive widespread use in electricity markets with dominant firms. In the longer term, they could desirably evolve in the direction of more comprehensive forward trading among suppliers and load. Such evolution will require changes in the regulatory structure such that both suppliers and load are given incentives or obligations to engage in forward trading — without this, an expansion of the supply offered in today's auctions would simply cause a collapse in prices and for reserve prices to become binding. But with such regulatory evolution, today's VPP auctions could provide a road map toward forward auctions where facilitating entry, developing wholesale markets, and reducing spot market power are all accomplished.

Acknowledgements

We gratefully acknowledge the support of National Science Foundation Grants SES-05-31254 and SES-09-24773. We are also grateful to Nils-Henrik von der Fehr and other participants in the IFN Workshop "Designing Electricity Auctions" in Stockholm for valuable comments. The authors were instrumental in setting the design of the EDF Generation Capacity Auctions, and Ausubel has subsequently worked as a consultant to EDF, Electrabel, E.ON, Endesa and Iberdrola in the implementation of their virtual power plant auctions. Consultants at IBM Global Business Services were instrumental in developing the product design for these VPP auctions. Further materials on VPP auctions are available on the websites of Power Auctions LLC (www.powerauctions.com) and Market Design Inc. (www.marketdesign.com), which served as subcontractors to IBM in the implementation of these auctions. Some of the methods described in this paper are proprietary; any discussion of methodology in this paper constitutes neither a recommendation nor a license to utilize any proprietary method.

References

Armstrong, Margaret, Galli, Alain, Kapoor, Arun., 2007. French VPP: Are buyers making a profit?, Working Paper, Cerna, Ecole des Mines de Paris, and Indian Institute of Technology, Delhi.

Ausubel, Lawrence M., 2004. An efficient ascending-bid auction for multiple objects. American Economic Review 94 (5), 1452–1475.

Ausubel, Lawrence M., Cramton, Peter, 2002. Demand Reduction and Inefficiency in Multi-Unit Auctions, University of Maryland Working Paper 96-07, revised July 2002.
Ausubel, Lawrence M., Cramton, Peter, 2004. Auctioning many divisible goods. Journal of the European Economic Association 2 (2–3), 480–493.
Ausubel, Lawrence M., Cramton, Peter. Using forward markets to improve electricity market design. Utilities Policy, in this issue.
Ausubel, Lawrence M., Cramton, Peter C., Jones, Wynne, P., 2002. System and Method for an Auction of Multiple Types of Items, International Patent Application No. PCT/US02/16937.
European Commission, 2001. COMP M. 1853-EDF/EnBW, Merger Procedure, Brussels.
Federico, Giulio, Vives, Xavier, Fabra, Natalia. 2008. Competition and Regulation in the Spanish Gas and Electricity Markets, Report of the Public-Private Sector Research Center, IESE Business School, University of Navarra.
Iberdrola, 2009. Iberdrola trades 62,000 GWh of power in Europe in 2008, more than triple the previous year's figure, press release, 16 January.
Schultz, Christian., 2005, Virtual Capacity and Competition, CESifo Working Paper No. 1487, University of Copenhagen.