The Hierarchical Clock Auction: A Modern Implementation of Hierarchical Package Bidding

Christina Aperjis[*] and Lawrence M. Ausubel[**]

23 August 2026

Very Preliminary

Abstract

Many previous attempts at designing computationally simple combinatorial auction formats have focused on auctions with prescribed hierarchical package bidding structures (Rothkopf, Pekeč and Harstad, 1998; Goeree and Holt, 2010). For example, the lowest level in the hierarchy could comprise the individual items, a middle level could consist of several non-overlapping packages of items, and the highest level could contain a single package comprising all items. With such a nested structure, solving the winner determination problem becomes trivial; at each level of the hierarchy, the highest bid for each package is compared to the sum of the highest bids for its component parts. Existing implementations are built on the simultaneous multiple round (SMR) auction format and inherit its weaknesses: bids from every earlier round remain live and may be revived many rounds later; minimum bids for items under a provisionally winning package are constructed by an arbitrary apportionment rule; regional bidders attempting to outbid a package face a free-rider problem; and the auction proceeds at the pace of an SMR auction.

This paper proposes a *Hierarchical Clock Auction* (HCA), which retains the economics of hierarchical package bidding while building on the ascending clock auction format used in modern spectrum, energy and resource auctions. Only current demand at current prices determines winners. In place of reviving earlier bids, the HCA reserves $K \geq 1$ of the $S \geq 2$ units in every region for bids placed at the regional level, so that at least one regional bid in every region is always at risk of winning; a national package is a separate product with its own price, which is prevented from running ahead of the regional prices and which rises on its own only when package bidders alone are competing. We give a complete specification and establish its basic properties. In a local–local–global (LLG) model, straightforward bidding is an ex post equilibrium of the clock stage and yields the efficient allocation, with Vickrey–Clarke–Groves payments when the package is sold; the outcome lies in the core and is a Walrasian equilibrium of the hierarchical economy. With several packages the format errs, if at all, only by selling too few. When a multi-region bidder faces a hard budget, the HCA reduces the budget to a simple exit rule, whereas under hierarchical package bidding dormant bids encumber the budget while purchasing nothing, and straightforward bidding becomes infeasible for the constrained bidder. Traditional hierarchical package bidding analyzed in the same model yields an allocation that depends on the order of bids and can expose regional bidders to winning licenses they do not want; the HCA achieves through its reserve and its pricing rule what bid revival and apportionment were designed to achieve.

---

[*] Power Auctions LLC, 1634 I St NW Ste 425, Washington, DC 20006. Email: aperjis@powerauctions.com

[**] Department of Economics, University of Maryland, College Park, MD 20742. Email: ausubel@econ.umd.edu

1. Introduction

Spectrum auctions routinely confront a tension between two kinds of bidders. A bidder with a geographically broad business plan values a nationwide footprint and faces an exposure problem if it must assemble that footprint license by license: it may win some of the licenses it needs and not others, paying for a partial footprint worth much less than the whole. A bidder with a local business plan values individual regions and is disadvantaged whenever a single all-or-nothing bid can sweep the region it wants. A fully general combinatorial auction or the combinatorial clock auction resolves the exposure problem, but at a cost: winner determination becomes computationally demanding, prices become opaque, and the threshold problem—the difficulty small bidders face in coordinating against a large package bid—becomes acute.

Hierarchical package bidding is the principal attempt to obtain the benefits of package bidding without paying that cost. Rothkopf, Pekeč and Harstad (1998) observed that when the permitted packages form a nested, tree-structured family, winner determination becomes computationally manageable: at each level of the hierarchy, the high bid on a package need only be compared with the sum of the high bids on its parts. Goeree and Holt (2010) built an auction format, Hierarchical Package Bidding (HPB), on this structure; it was tested experimentally by Brunner, Goeree, Holt and Ledyard (2010) and by Kagel, Lien and Milgrom (2010), and it was used by the Federal Communications Commission in the 2008 auction of 700 MHz licenses (FCC, 2007), whose rules implement the format in full. HPB occupies an attractive middle ground: winner determination is trivial and transparent, the format is a modest extension of the familiar simultaneous multiple-round (SMR) auction, and the hierarchy can be designed to match the footprints broad bidders are likely to want.

The rules of the 2008 auction also reveal the format's weaknesses. They arise because HPB is built on an SMR auction—a format in which a bid is a standing offer rather than an expression of demand at a current price. We identify four issues:

1. **Bid revival.** When a package bid is the provisionally winning bid (PWB), the high bids on its constituent items are not PWBs, but they remain live. If the package is later displaced—by a higher package bid, or because rising bids on other constituents push the item sum above the package bid—the dormant item bids become PWBs again, possibly many rounds after they were placed and possibly for a bidder who has long since stopped bidding. Revival is what keeps item bids relevant under HPB, but it does so by imposing an open-ended liability on item bidders and by making the final allocation depend on stale information. A bidder can be surprised, many rounds later, to win items for which it had stopped bidding. This is especially an issue for budget-constrained bidders. Depending on the activity rule, a bidder may meanwhile be constrained by bids that are not provisionally winning.

2. **Current price estimates.** While a package is the PWB, the minimum acceptable bid for each constituent item is derived from a "current price estimate," constructed by scaling the item bids up, in proportion to bidding units, until they sum to the package bid. The apportionment is arbitrary; it can make the minimum bid for a thin region jump far above

any bid ever placed there; and it is the channel through which the threshold problem operates under HPB.

3. **The free-rider problem.** Regional bidders attempting to outbid a package must collectively cover the shortfall between the package bid and the item sum. Each would prefer that the others cover it, and an item bidder whose bid is dormant beneath a provisionally winning package can wait, at no cost, for the others to act. Regional bidders are thereby placed at a structural disadvantage relative to the package bidder.

4. **Pace.** Like any SMR auction, HPB raises the price of an item only when a new bid is placed on it. With one hundred regions and several blocks per region, the auction can run for a very large number of rounds.

Ascending clock auctions, which have become the international standard format for spectrum auctions[1] and are also used in energy[2] and other resource auctions[3], address the fourth issue directly and the first three issues indirectly. In a clock auction the auctioneer names prices and bidders respond with quantities. A bid is an expression of demand at the current price; there are no provisionally winning bids and nothing to revive; the prices of all units of a product rise together; and the auction ends when demand no longer exceeds supply. The difficulty remains that the standard clock design has no packages: a bidder that wants a nationwide footprint must hold demand in every region separately, and is exposed in exactly the way HPB was designed to prevent.

This paper proposes an ascending clock auction—the *Hierarchical Clock Auction* (HCA)—that combines the two. The design is deliberately cautious: every element of the standard clock format is retained wherever possible.

- **A single regional product per region.** Each region's S generic blocks constitute one product with one price, exactly as in a standard clock auction. Any bidder may bid for them at the regional level.

- **A national package product with its own price.** The national package comprises one block in every region, demanded on an all-or-nothing basis. Its price normally equals the sum of the regional prices and is never permitted to run ahead of them while regional competition is driving prices; it rises on its own only when package demand is excessive and no region has regional excess demand, and any premium so created is absorbed as the regional prices subsequently rise.

---

[1] The format—a clock stage with generic blocks, intra-round bidding, sequential bid processing under a no-excess-supply rule and an eligibility-based activity rule, followed by an assignment stage—has been used, among others, in the United States for almost every major spectrum auction since the 2016–17 broadcast incentive auction, in Canada (3500 MHz, 2021; 3800 MHz, 2023) and in Australia (3.6 GHz, 2018; 26 GHz, 2021). Its rules are set out in the published procedures of those auctions (e.g., FCC, 2019; ISED, 2020).

[2] See Ausubel and Cramton (2010).

[3] See BOEM (2024).

- **Reserved units in every region.** $K \geq 1$ of the S blocks in each region are reserved for sale at the regional level, so that at most $S - K$ blocks in any region can be won through the package. The clearing condition of the clock stage is therefore two-fold: aggregate demand in every region (regional-level plus package demand) is at most S *and* package demand is at most $S - K$. This is precisely the treatment of set-aside spectrum in recent clock auctions (ISED, 2020; Section 7.2), in which set-aside-eligible bidders may win any block but set-aside-ineligible bidders may collectively win no more than the open blocks—adapted from a restriction on bidder classes to a restriction on bid types.

The reserved units are the substitute for bid revival. Under HPB, a regional bidder's protection against being swept by a package is the possibility that its dormant bid will be revived. Under the HCA, its protection is structural: package demand cannot fill the last K units in any region, so demand expressed by regional-level bidders at the current price always determines who wins at least K units in every region. No bid is ever revived; every winning demand is an expression made at the final price; and a regional bidder's liability is never more than its current demand at the current price. Because the package has its own price, current price estimates are unnecessary: whenever regional competition moves the package price, that price is simply the sum of the regional prices.

The subtlest of the new rules governs the meeting of all-or-nothing package demands with the "no excess supply" rule of the standard format. Section 4.7 resolves it with a *package exit rule*: a package reduction is applied, in its entirety, if and only if, at the moment it is processed, the package price is being driven up—because package demand exceeds $S - K$, or because some region has regional excess demand. A package bidder held to a package it asked to relinquish therefore never pays more than the price at which it asked to relinquish it. The cost of the rule is that a package exit can leave undersell—units in regions where the departing bidder was the marginal demander. The clock stage may close with undersell; an optional supplementary stage (Section 5) places undersold units in bidders' hands before the assignment stage.

Our principal results concern straightforward bidding in the tradition of Milgrom (2000) within a local–local–global (LLG) environment in the tradition of Krishna and Rosenthal (1996): regional bidders with unit demands in a single region, and global bidders whose value is for the national package. Theorem 1 shows that straightforward bidding in the HCA yields the efficient allocation, selling the package exactly when the global bidder's value exceeds the displacement value—the sum across regions of the values of the regional bidders the package displaces—with Vickrey–Clarke–Groves payments for every winner when the package is sold. Theorem 2 shows that the outcome lies in the core, that every payment is at least the Vickrey payment, and that the final prices constitute a Walrasian equilibrium of an economy in which the package and the regional blocks are distinct goods. Theorem 3 extends the analysis to several packages: the HCA never sells too many; its only possible error is to sell too few, and the welfare loss is bounded by the spread of regional values. Proposition 3 shows that straightforward bidding is an ex post equilibrium of the clock stage. Proposition 4 introduces a budget-constrained multi-region bidder and shows that the HCA reduces the budget to a simple, belief-free exit rule, whereas under HPB straightforward bidding is incompatible with budget feasibility, because dormant bids encumber the budget while purchasing nothing. Finally, Proposition 2 places HPB in the same environment, under the same

behavioral assumption: its allocation depends on the order in which bids are placed; its revival of dormant bids can hand a regional bidder a second license it does not want while excluding a bidder who would pay the price; and the increments it imposes on regional bidders are apportioned shares of *package* increments, at the scale of the package rather than of the region. We are careful about what this comparison does and does not establish: the inefficiencies of HPB under straightforward bidding are increment effects, and the threshold problem proper is a question about equilibrium behavior that we leave open.

The paper is related to several elements of the auction literature. The simultaneous multiple round (SMR) auction, with its exposure and threshold problems, is analyzed by Milgrom (2000, 2004). Efficient ascending clock auctions are developed in Ausubel (2004, 2006). The combinatorial clock auction (CCA), which avoids the exposure problem and has also been used in spectrum auctions, is proposed in Ausubel, Cramton and Milgrom (2006). The LLG model is due to Krishna and Rosenthal (1996) and is the standard environment for studying core-selecting auctions; see Day and Milgrom (2008) and, for the incomplete-information case, Ausubel and Baranov (2020). Budget-constrained bidding in auctions is analyzed by Che and Gale (1998), while Bulow, Levin and Milgrom (2009) document the central role of budget constraints in the US 700 MHz auction (the one auction in which HPB has been used in the field). Finally, demand reduction in uniform-price auctions is analyzed by Ausubel, Cramton, Pycia, Rostek and Weretka (2014).

The remainder of the paper is organized as follows. Section 2 sets out the environment and the HPB benchmark. Section 3 gives an overview of the HCA. Section 4 is the complete specification of the clock stage; Section 5 describes the supplementary and assignment stages. Section 6 contains the results. Section 7 discusses the design rationale in the light of the results, Section 8 presents variants and extensions, including a sketch of an intermediate tier of regional packages, and Section 9 concludes. Appendix A collects the proofs.

2. The Environment and the HPB Benchmark

2.1 Licenses, regions and blocks

There is a set $R = \{1, \dots, j, \dots, J\}$ of regions, $|R| = J$. Each region has $S \geq 2$ blocks, labeled $A, B, \dots, X$, sufficiently similar in bandwidth, encumbrance and technical characteristics that bidders treat them as close substitutes, up to considerations—contiguity within a region, the same block across regions, harmonic interference with other spectrum holdings—naturally handled in an assignment stage. We therefore treat the S blocks in a region as S units of a generic product, and call S the supply of the regional product. Of these, K units ($1 \leq K \leq S - 1$) are reserved for sale at the regional level (Section 3). Our main description assumes a common S and K across regions; Section 8.1 covers region-specific values.

Each region carries bidding units $b_j > 0$, used to measure activity, and a minimum opening bid $m_j > 0$ per block. A spectrum cap may limit the blocks, including prior holdings, that a bidder may hold in a region; $c_{i,j}$ denotes the maximum number of blocks bidder i may win in region j.

2.2 The package

In the base case there is a single package, the national package N: one block in every region, with bidding units $b_0 = \sum_{j \in R} b_j$. A bidder that wants two blocks nationwide demands two units of N. The extension to an intermediate tier—for example, a partition of R into four or six regional packages—is sketched in Section 8.2. The clearing rules of the base case have been fully worked out; those of the extension have not.

2.3 Bidders and preferences

Bidders $i \in I$ are of two broad kinds, although the format does not classify them. *Regional bidders* value blocks in particular regions, with little or no complementarity across regions. *Package bidders* value a nationwide footprint and suffer a substantial value reduction from missing any components. The exposure problem is the risk, borne by a package bidder, of winning an unprofitable partial footprint; the threshold problem is the difficulty, faced by a group of regional bidders, of collectively outbidding a package bidder when each would prefer that the others contribute more.

Bidders have initial eligibility, measured in bidding units and determined by upfront payments; may qualify for bidding credits; and may be subject to spectrum caps. All are handled exactly as in the standard clock format.

2.4 The HPB benchmark

The HPB format, as implemented in the 2008 U.S. 700 MHz auction (Goeree and Holt, 2010; FCC, 2007), operates as follows. Bidders bid in rounds on individual items or on pre-defined packages that partition the items; a bidder may not bid on an item and on a package containing it in the same round. After each round the system identifies the high bid on each item and each package, treating every item as carrying a seller bid just below its minimum opening bid. If the sum of the high bids on a package's items is at least the high bid on the package, the item bids become the provisionally winning bids (PWBs); otherwise the package bid is the PWB. Item bids that are not PWBs remain live. Consequently, an increase in the high bid on one or more items can lift the item sum above the package bid and thereby "revive" dormant bids on other items—placed in earlier rounds, perhaps by bidders no longer active—back into being PWBs; and the comparison between package bid and item sum can go back and forth repeatedly. The minimum acceptable bid for each item is computed from a current price estimate (CPE): when a package bid is the PWB, each item's CPE adds to its highest bid to date a share of the shortfall proportional to its bidding units. The auction ends after a round with no new bids, no withdrawals and no proactive waivers.

The format has genuine virtues: trivial and transparent winner determination, no exposure for the package bidder, and a guarantee that an item bid, once placed, can win if the package is displaced. But the revival of earlier bids is also its central weakness. The item bidder has an open-ended commitment; the final allocation can rest on bids placed at prices bearing no relation to the final prices of neighboring items; a bidder that has left the auction can be handed licenses it no longer wants and (due to other winnings) it no longer can afford; and the CPE apportionment, rather than any bidder's valuation, determines how hard it is for a region's bidders to outbid a package. The HCA is designed to keep the virtues but remove its main weakness.

3. Overview of the Hierarchical Clock Auction

The HCA consists of a clock stage, an optional supplementary stage, and an assignment stage. The clock stage determines how many generic blocks each bidder wins in each region—at the regional level or through the national package—and the price of each; the supplementary stage, if used, disposes of any undersell; the assignment stage determines which specific blocks each winner receives. We summarize each element here; Sections 4 and 5 give the complete specification.

Products. There are $J + 1$ products: the regional product of each region j (generic blocks), with supply S; and the national package N (comprising one generic block in every region), demanded on an all-or-nothing basis, with supply at most $S - K$. Any bidder may demand any product.

The reserve. At most $S - K$ units of the package are permitted to be sold and so, implicitly, K out of the S units in each region are reserved for sale at the regional level. The reserve restricts bid types, not bidders: a bidder demanding the package may also demand regional units, subject to a per-bidder limit of S blocks per region across both kinds of demand and of $S - K$ package units in total.

Prices and eligibility. Each regional product has a clock, and the package has a clock of its own. The package price normally equals the sum of the regional prices: while at least one region has regional excess demand, the package clock price is the sum of the regional clock prices, so the package price cannot run ahead of the regional prices. If a point is reached when no region has regional excess demand, the package clock price is instead a specified percentage above the previous round's posted package price, so that competition among package bidders alone can be resolved without lifting the prices faced by regional bidders. The difference between the package price and the sum of the regional prices—the "package premium"—is always non-negative, arises only from competition among package bidders, and is absorbed as regional prices subsequently rise. Bidders hold separate eligibility for package and for regional-level bidding.

Rounds and bids. The auction proceeds in timed rounds. In each round after the first, every product has a start-of-round price (the previous round's posted price) and a clock price. A bidder states, for each product, the quantity it demands at the clock price, and may submit intra-round bids changing its demand at price points between the two. Package demand and regional-level demand are distinct and are never converted into one another: no bid, and no processing rule, turns a package unit into regional units or the reverse.

Bid processing. Bids are processed sequentially, using the queue-and-retest procedure of the standard clock auction format, in three steps: bids to maintain demand first; then bids to reduce package demand, all submitted at the start-of-round price; and then all other bids to change demand, in order of their price point with a random tie-break. Increases are applied subject to eligibility, the per-bidder limits and spectrum caps. Regional reductions are applied to the maximum extent that does not create or increase excess supply, with package demand counted at no more than $S - K$. Package reductions are applied in their entirety if, at the moment of processing, package demand exceeds $S - K$ or some region has regional excess demand; otherwise they are held at the start-of-round price.

Clearing and the next round. The clock stage ends after the first round in which aggregate demand in every region is at most S and package demand is at most $S - K$. Winners are the bidders

with processed demand in the final round; regional winners pay the final regional posted prices, while package winners pay the final package posted price. Undersell is possible and is carried into the supplementary stage.

4. Clock Stage: Complete Specification

We now specify the clock stage completely. The reader interested primarily in the results may proceed from Section 3 to Section 6, which refers back to the specification at each step.

4.1 Notation

The set of regions is $R = \{1, \dots, J\}$, with typical element j, and the national package is denoted by 0. Each region has a supply of S generic blocks, of which K are reserved for sale at the regional level, $1 \leq K \leq S - 1$. The package supply $S_0 = S - K$ is the maximum number of package units that can be sold. The j^{th} regional product carries b_j bidding units, and the package carries $b_0 = \sum_j b_j$ bidding units.

Bidder i's processed demands are denoted $x_{i,j}$ regional-level units in region j and $x_{i,0}$ package units. Aggregating over bidders, $D_j = \sum_i x_{i,j}$ is regional-level demand in j, $D_0 = \sum_i x_{i,0}$ is package demand, $AD_j = D_j + D_0$ is aggregate demand in region j, and $\overline{AD}_j = D_j + \min\{D_0, S_0\}$ is *capped* aggregate demand, in which package demand is counted at no more than the package supply.

In round t, the regional product of j has posted price p_j^t and clock price P_j^t; the package has posted price p_0^t and clock price P_0^t; and the *package premium* is $\psi^t = p_0^t - \sum_j p_j^t$. Finally, bidder i's regional and package eligibilities in round t are denoted E_i^t and $\hat{E}_i^t$, respectively, and its processed regional and package activities are denoted a_i^t and $\hat{a}_i^t$, respectively, all measured in bidding units.

4.2 The reserve and per-bidder limits

Package demand can be filled from at most S_0 units in any region, so at least K units in every region go—if they are won at all—to regional-level demand. The reserve K is a design parameter (Section 8.1 covers region-specific values).

Two per-bidder limits are enforced at submission: $x_{i,0} \leq S_0$, and total demand in any region—package plus regional-level units—of at most the supply, $x_{i,0} + x_{i,j} \leq S$ for every j. Spectrum caps apply on top.

The reserve restricts *bids*, not *bidders*. A bidder demanding the package may also demand regional units anywhere, and a bidder with nationwide interests may bid entirely at the regional level. What the reserve excludes is package demand filling every unit in a region.

No conversion. Package demand and regional-level demand are distinct products, and neither is ever converted into the other: no bid type exchanges package units for regional units or the reverse, a released package unit is simply released (it does not become regional-level demand anywhere), and regional-level units are never aggregated into a package. A bidder pursuing both routes holds

demand—and separate eligibility (Section 4.5)—for both, its bids for the two products processed independently under each product's own rules.

4.3 Excess demand and excess supply

The clearing condition of the clock stage has two parts, and excess demand is defined for each.

Package excess demand. The package has excess demand when $D_0 > S_0$: more package units are demanded than can be won.

Regional excess demand. The regional product of region j has excess demand when

$$\overline{AD}_j = D_j + \min\{D_0, S_0\} > S,$$

that is, when regional-level demand exceeds the units left over after package demand, counted at no more than the package supply, has been accommodated. When $D_0 \leq S_0$ this is simply $AD_j > S$; when $D_0 > S_0$ it is $D_j > K$. With $S = 2$, $K = 1$ and at least one package bidder, region j has excess demand exactly when more than one regional-level unit is demanded there.

The regional product has **excess supply** when $\overline{AD}_j < S$. (Because package demand above the package supply cannot be filled, it is the capped aggregate that determines whether units are unfilled.)

The capped test is the test used for set-aside products in clock auctions with set-asides (ISED, 2020; Section 7.2), with regional-level demand as set-aside demand, the K reserved units as the set-aside supply, and package demand as open demand. In both formats, a regional (set-aside) bidder counts as competing only for a unit it could lose to someone else.

Clearing. The clock stage clears when no regional product has excess demand and the package has no excess demand:

$$AD_j \leq S \quad \text{for every } j \in R \qquad \text{and} \qquad D_0 \leq S_0.$$

(Given $D_0 \leq S_0$, the capped and uncapped aggregates coincide.)

4.4 Prices

Round 1. The start-of-round price and the clock price of every regional product equal the minimum opening bid m_j; the package's equal $p_0^1 = \sum_j m_j$. In Round 1 a bidder simply states the quantity of each product it demands at the opening prices.

Regional clock prices, rounds $t \geq 2$. The start-of-round price of the regional product of j is the previous round's posted price p_j^{t-1}, and its clock price is a specified percentage above it,

$$P_j^t = (1 + \gamma^t)\, p_j^{t-1},$$

where γ^t is the round's increment percentage (initially 5–10 percent; a dollar cap may also be imposed). The clock price is set for every regional product, whether or not it has excess demand; a product without excess demand posts at its start-of-round price again unless new demand arrives.

Package clock price, rounds $t \geq 2$. The start-of-round price of the package is p_0^{t-1}. Its clock price depends on the state at the end of round $t-1$:

- *Case (a): some region has regional excess demand.* The package clock price is the sum of the regional clock prices, unless the package already stands above that sum:

$$P_0^t = \max\left\{p_0^{t-1},\ \sum_{j \in R} P_j^t\right\}.$$

- With no premium ($\psi^{t-1} = 0$) this is exactly $\sum_j P_j^t$; with a premium, the regional clock prices catch up with the package price before it moves again—the premium is absorbed, not carried. Either way the package price cannot get ahead of the regional prices while regional competition is moving them.
- *Case (b): no region has regional excess demand.* The package clock price is a specified percentage above the previous posted package price,

$$P_0^t = (1 + \gamma^t)\, p_0^{t-1}.$$

- This is the only circumstance in which the package price moves independently of the regional prices; it is what allows competition among package bidders alone to be resolved without lifting the prices faced by regional bidders.

(Absent a premium the two cases coincide, since $\sum_j (1+\gamma^t)\, p_j^{t-1} = (1+\gamma^t) p_0^{t-1}$; they differ in whether an existing premium is absorbed by the regional prices, as in case (a), or scaled with the package price, as in case (b). In both, $P_0^t \geq p_0^{t-1}$ and $P_0^t \geq \sum_j P_j^t$.)

Price points. A bid at price p for a product with start-of-round price p^s and clock price P has price point $\theta = (p - p^s)/(P - p^s)$: $\theta = 0$ is the start-of-round price, $\theta = 1$ the clock price. Regional bids use the regional product's range; package increases use the package's own range, a package bid at θ being at price $p_0^{t-1} + \theta(P_0^t - p_0^{t-1})$. With no premium, in case (a), the package price at θ equals the sum of the regional prices at the same θ. Package reductions are submitted at the start-of-round price only (Section 4.6).

The premium. Section 4.8 shows that $p_0^t \geq \sum_j p_j^t$ in every round, so the premium ψ^t is non-negative. It starts at zero, rises only in a round in which the package ends with excess demand and posts at its clock price while some regional products post below theirs, and otherwise falls as regional prices rise, until absorbed (Section 6.1). The premium is the amount by which the package route to a block costs more than a regional-level bid—the HCA counterpart of the gap between open and set-aside prices in clock auctions with set-aside blocks (ISED, 2020; Section 7.2)—and, as there, it arises only from competition among the bidders who cannot use the reserved route.

4.5 Eligibility, activity and caps

Separate eligibility. Each bidder has two eligibilities, each in bidding units and determined by its upfront payment: a *regional eligibility* E_i^1 governing regional-level demand and a *package eligibility* $\hat{E}_i^1$ governing package demand. Each can be used only for bids of its own kind, and neither is ever transferred to the other; a bidder intending only one route acquires eligibility for that route alone.

Separate eligibility is the activity-rule counterpart of the no-conversion rule of Section 4.2: activity freed by releasing a package unit cannot take up regional-level demand, nor the reverse.

Activity. A bidder's regional activity is the bidding units of its regional-level demand, a unit in region j counting b_j; its package activity is the bidding units of its package demand, a unit counting b_0. Its requested activity of each kind is the activity of its requested demand at the clock prices; its processed activities after the round are

$$a_i^t = \sum_{j \in R} x_{i,j}\ b_j, \qquad \hat{a}_i^t = x_{i,0}\, b_0.$$

Activity rule. The activity requirement in round t is a percentage $\alpha^t \in [90\%, 100\%]$, initially 95 percent and raised to 100 percent in a later phase. It applies to each eligibility separately: if $a_i^t \geq \alpha^t E_i^t$, regional eligibility is unchanged; otherwise $E_i^{t+1} = a_i^t / \alpha^t$; likewise for package eligibility. Neither eligibility ever increases. There are no activity-rule waivers: as in the standard design, the price-point mechanism lets a bidder express demand at every price in the round. The "grace round" formula of some implementations (ISED, 2020)—eligibility computed from the larger of processed activity and the smaller of the previous round's processed and the current round's submitted activity—may be substituted for either eligibility without affecting anything else. Because package units are indivisible multiples of b_0, package eligibility is in practice a whole number of package units, and a bidder that releases a unit loses that unit's eligibility at the next round.

Activity upper limit. Optionally, as in some implementations, a bidder may *submit* regional-level bids whose requested activity exceeds its regional eligibility by up to a stated percentage (e.g., 120 percent), so that a bidder whose reductions are held in the queue can still place offsetting increases; the system never *applies* bids carrying processed regional activity above eligibility. No such limit is needed for the package, since package reductions are all-or-nothing and package increases are checked against package eligibility alone.

Spectrum caps and per-bidder limits. A submission is rejected if at any price point it would violate $x_{i,0} \leq S_0$, $x_{i,0} + x_{i,j} \leq S$ for any j, or a spectrum cap $c_{i,j}$ (each package unit counting as one block in every region).

Bidding credits. Clock-stage prices are gross. Net commitments apply the bidder's credit percentage subject to any caps the seller imposes, the package commitment $x_{i,0}p_0^t$ being apportioned to regions in proportion to regional posted prices for any region-specific cap.

4.6 Bids

Bid types. In Round 1 a bidder states a quantity for each product at the opening price. In every later round, for each product a bidder may:

1. *Maintain demand* at the clock price. Always applied.
2. *Reduce regional-level demand* to a quantity below its processed demand at a price point θ: the bidder is willing to buy its previous quantity at all prices below θ, any quantity between the two at θ, and the requested quantity above θ. Several reductions at increasing price points with decreasing quantities may be submitted.

3. *Reduce package demand* to a quantity below its processed demand, **at the start-of-round price only**: the bidder is willing to buy its previous package quantity at the start-of-round price and the reduced quantity at any higher price in the round. A package reduction cannot be conditioned on a price point within the round.
4. *Increase demand* for a regional product or for the package to a quantity above its processed demand at a price point θ: the bidder is willing to buy any quantity between its previous demand and the requested quantity at all prices in the round. Several increases at increasing price points with increasing quantities may be submitted.
5. *Switch demand between regional products*: a linked pair consisting of a reduction in one or more regional products and an increase in one or more regional products at the same price point, the increase to be applied only to the extent that the reduction is applied, as in the standard design. A bidder's bids for a product within a round must be one-directional.

Why package reductions are at the start-of-round price. The restriction simplifies bid processing considerably (Section 4.7): package reductions never interleave with regional bids, the package's posted price is determined by at most two rules, and a held package bidder pays exactly the price at which it asked to leave. Its cost is that a package bidder's exit decision is made round by round rather than at an arbitrary price within the round: a bidder maintaining the package accepts that its price may rise as far as that round's package clock price. The increment percentage bounds the rise, so the seller controls the size of this commitment.

No switches across the hierarchy. A switch may not link the package with regional products. A bidder may, in the same round, reduce its package demand and increase its regional-level demand, or the reverse, but these are independent bids for independent products: each is processed on its own terms, neither is conditioned on the other, and each increase draws on the eligibility of its own kind. In particular, a bidder that releases a package unit cannot take up regional-level units with the freed package activity, then or later; it can bid regionally only with regional eligibility it already holds. The absence of cross-hierarchy switches is deliberate: the two routes to a block are meant to be chosen, not traded between (Section 7.2).

Missing bids. A bidder with positive processed demand for a product that submits no bid for it is deemed to have bid zero at the start-of-round price; for the package, the deemed bid is a reduction processed under the exit rule.

Validation. The system rejects a submission violating the per-bidder limits or caps at any price point, or whose requested activity exceeds the applicable eligibility (or activity upper limit). Bids may be revised until the round closes.

4.7 Bid processing

Bid processing after each round determines processed demands and posted prices. It follows the published procedure of the standard clock format (e.g., FCC, 2019; ISED, 2020) except in three respects: the placement of package reductions in the processing order, their treatment, and the use of the capped aggregate.

Order of processing. Bids are processed in three steps.

1. *Bids to maintain demand* are applied first and unconditionally.
2. *Bids to reduce package demand*, all of which are at the start-of-round price, are processed next, in an order given by a bid-specific pseudorandom number drawn once for the round. Each is tested under the package exit rule below; a package reduction not applied in its entirety is placed in the queue.
3. *All other bids to change demand*—regional reductions, regional and package increases, and switches among regional products—are then processed in order of price point from lowest to highest and, within a price point, by pseudorandom number.

The system maintains a **queue** of bids not yet applied in their entirety. Whenever any bid is applied, in whole or in part, every queued bid is re-tested, in order; re-testing iterates until nothing more can be applied, and only then is the next bid processed. When every bid has been processed and nothing in the queue can be applied, the queue is discarded.

Regional increases. A bid to increase regional-level demand in region j is applied one unit at a time, to the maximum extent that keeps the bidder's processed regional activity within its regional eligibility, its holdings within $x_{i,0} + x_{i,j} \leq S$, and its holdings within its spectrum cap.

Regional reductions (no excess supply rule). A bid to reduce regional-level demand in region j is applied one unit at a time while $\overline{AD}_j > S$, evaluated at the moment of processing with all demands applied so far. A reduction is never applied into excess supply, and not at all if the product already has excess supply. This is the standard rule with the capped aggregate in place of the plain one; with package demand counted at no more than S_0, a regional bidder is held only when its unit would otherwise go unfilled.

Package increases. A bid to increase package demand is applied one unit at a time, to the maximum extent that keeps processed package activity within package eligibility, $x_{i,0} \leq S_0$, $x_{i,0} + x_{i,j} \leq S$ in every region, and holdings within caps. An added package unit raises D_0 by one and, if $D_0 \leq S_0$ afterwards, raises $\overline{AD}_j$ by one in every region.

Package reductions (package exit rule). A bid to reduce package demand from z to $z' < z$ units is processed one unit at a time. A package unit is **released** if and only if, at the moment of processing, the package price is being driven up—that is, if either

$$D_0 > S_0 \qquad \text{or} \qquad \overline{AD}_j > S \text{ for some } j \in R.$$

The release is never partial. It lowers D_0 by one; if $D_0 \leq S_0$ *before* the release—so that the released unit was counting toward capped aggregate demand—it also lowers $\overline{AD}_j$ by one in *every* region, and it is permitted to create or increase excess supply in regions other than the one(s) with excess demand. A release applied while $D_0 > S_0$ relieves package excess demand and leaves every $\overline{AD}_j$ unchanged. If neither condition holds, the unit is not released, the bid remains in the queue, and it is re-tested whenever any other bid is applied. A package unit not released by the end of processing is **held**: the bidder's processed package demand remains at its previous level, and—because the reduction was at the start-of-round price and nothing drove the package price up—the bidder pays exactly the start-of-round package price (Proposition 1(c)).

Three remarks explain the rule and its placement. First, the two conditions are exactly the circumstances in which the package's posted price can rise (Section 4.8): package excess demand lifts it directly, regional excess demand through the sum. A package bidder may leave whenever, and only whenever, someone else's demand is driving up the price of its package. Second, because package reductions are processed before any other bid to change demand, they are tested against the state at the close of the previous round—a state that failed the clearing condition, or the round would not be taking place. The first package reduction processed in any round is therefore always applied. Later ones may be held: once package demand has fallen to the package supply and the releases have removed the regional excess demand, nothing is driving the package price up and the remaining package bidders are held at the start-of-round price. Third, a held package reduction stays in the queue, and any later bid that creates excess demand—a regional increase that makes a region hot, or a package increase pushing package demand above the package supply—triggers its release before that excess demand can move the package price. Section 7.3 discusses why the alternatives—partial application region by region, or unconditional release—are inferior.

Switch bids between regional products. The increase leg of a switch is applied only to the extent that the reduction leg has been applied, and is then subject to the checks above. No switch involves the package.

Consequences for a held bidder. A bidder whose reduction is held—regional, by the no-excess-supply rule, or package, by the exit rule—retains its previous processed demand and pays the posted price. Proposition 1(c) shows that a held regional bidder pays no more than the price at which it asked to reduce, and a held package bidder exactly the start-of-round package price.

4.8 Posted prices

Regional products. Once processed demands are determined, the posted price of the regional product of j is set as follows:

- if the product has excess demand, $\overline{AD}_j > S$, the posted price is the clock price P_j^t;
- otherwise, if at least one regional reduction in j was applied in the round, entirely or partially, the posted price is the highest bid price among the regional reductions applied—the price at which a reduction brought demand down to supply;
- otherwise, the posted price is the start-of-round price p_j^{t-1}.

This is the standard posted-price rule with one clarification: a package release also reduces capped aggregate demand in region j, but since every package reduction is at the start-of-round price, a release can never raise a regional posted price, and it is disregarded in setting posted prices.

The package. The posted price of the package is set as follows:

- if the package has excess demand, $D_0 > S_0$, the posted price is the package clock price P_0^t;
- otherwise, the posted price is the sum of the regional posted prices, unless the package already stands above that sum: $p_0^t = \max\{p_0^{t-1}, \sum_j p_j^t\}$.

The second rule is the sum rule: the package price follows the regional prices, and any premium is absorbed as they rise. The standard rule that a product without excess demand posts at the price of

the reduction that brought demand down to supply is subsumed, every package reduction being at the start-of-round price, which the rule already includes. With no premium, these rules coincide with the standard posted-price rule applied product by product.

The premium is non-negative and is absorbed. The second rule gives $p_0^t \geq \sum_j p_j^t$ directly, and $P_0^t \geq \sum_j P_j^t \geq \sum_j p_j^t$ in both cases of Section 4.4, so the premium ψ^t is non-negative in every round. It increases only when the package posts at its clock price while some regional products post below theirs—only, that is, when package competition is lifting the package price—and in every other round it is $\max\{\psi^{t-1} - \sum_j \left(p_j^t - p_j^{t-1}\right), 0\}$: regional price increases eat into it until it is gone. Sections 6.3 and 6.8 show that this absorption is what allows regional bidders to displace a package bidder only when their aggregate willingness to pay exceeds its own.

Releases into excess supply. When a package release leaves region j with excess supply, the stranded unit remains on offer at the regional posted price: the highest price at which a regional reduction brought demand down to supply before the release, or the start-of-round price.

4.9 Eligibility and prices for the next round

After processing, each bidder's processed activities and next-round eligibilities are computed as in Section 4.5, and the next round's start-of-round and clock prices as in Section 4.4, with the case distinction for the package clock price made on the regional excess-demand status at the end of the round just processed. Bidders are informed of their processed demands, processed activities, eligibilities, and gross and net commitments at posted prices.

4.10 Stopping rule, winners and undersell

The clock stage ends after the first round T in which the clearing condition holds: $AD_j \leq S$ for every region and $D_0 \leq S_0$. If an aggregate reserve price applies and is not met, the auction ends with no licenses assigned.

Each bidder wins the generic blocks of its processed demand after round T: $x_{i,j}$ blocks in region j at p_j^T, and $x_{i,0}$ package units, each one block in every region, at p_0^T. Gross payments are commitments at final posted prices; net payments apply bidding credits.

Undersell. Regions with $AD_j < S$ at the close have $S - AD_j$ unsold blocks. Undersell arises in two ways: a reserved unit for which no regional-level demand was ever expressed at its posted price, and a non-reserved unit stranded when a package release removed its marginal demand. Neither prevents the clock stage from ending; the optional supplementary stage of Section 5.1 then offers the undersold units before the assignment stage. Under HPB, the corresponding units revert to the seller's bid in exactly the same circumstances.

4.11 Information policy

After each round the system publishes, for every region, the supply, D_j, AD_j, the posted price and the next clock price; and, for the package, D_0, the package supply, the posted price, the next clock price and the current premium. Separate disclosure of regional-level and package demand is the

counterpart of reporting set-aside and open demand separately: every bidder can see whether a region's price is moving because of regional or package competition. Under limited disclosure, bidder identities are withheld until the clock stage ends; under full disclosure, processed demands are published each round. There are no provisionally winning bids to report and, once a request to reduce quantity has been applied, no possibility of a past bid converting to a winning bid.

5. Supplementary Stage and Assignment Stage

5.1 Supplementary stage (optional)

If the clock stage closes with undersell, the seller may conduct a supplementary stage before the assignment stage, whose sole purpose is to place undersold blocks in bidders' hands. It does not reopen the clock stage, cannot change clock-stage winnings or prices, and is announced in advance in the auction procedures.

Offer. For each region j with undersell, $u_j = S - AD_j$ generic blocks are offered, at an opening price equal to the final clock-stage posted price p_j^T. (The seller may instead announce that the supplementary stage opens at the clock-stage opening price, or one increment above p_j^T; Sections 6.5 and 6.8 show what turns on the choice.) Blocks are offered at the regional level only; there is no package in the supplementary stage.

Participation and eligibility. A bidder's supplementary-stage eligibility is its initial regional eligibility E_i^1: the stage continues the regional route, and a bidder that acquired eligibility only for the package cannot bid in it. (At the seller's option, bidders may instead establish fresh eligibility by an additional deposit; Section 8.6 gives the safeguard then appropriate.) Clock-stage plus supplementary-stage winnings may not exceed S in any region, and spectrum caps continue to apply.

Format. The supplementary stage is itself an ascending clock auction over the undersold blocks, with opening prices p_j^T, the rules of Section 4 applied with package demand absent (aggregate demand for region j is simply D_j, supply is u_j), and the standard stopping rule: it ends after the first round in which no region has excess demand. Two versions differ in whether demand can move between regions.

- *Parallel clock auctions (no switching).* Each region with undersell is a separate single-product clock auction, run simultaneously but independently. In round 1 a bidder states its demand in each region at the opening price; thereafter it may only maintain or reduce, reductions being subject to the no-excess-supply rule. With no increases and no switches, demand can never move between regions, and a bidder's eligibility is in effect its round-1 demand region by region. A region's auction ends when its demand no longer exceeds its undersell; the stage ends when every region's has.
- *Clock auction with switching.* The undersold blocks are offered in a single multi-product clock auction in which supplementary-stage eligibility is a common pool of bidding units. Bidders may maintain, reduce, increase and switch demand among regions at price points,

subject to eligibility, the no-excess-supply rule and the activity rule, exactly as in the clock stage. The stage ends after the first round with no excess demand anywhere.

Under either version, any undersold block for which demand reaches supply is sold, by the no-excess-supply rule; blocks not demanded at their opening price remain unsold. Winners pay the stage's final posted prices. The parallel version is simpler and lets a bidder interested in a few regions ignore the others; the switching version lets bidders substitute among regions as prices move, at the cost of a full multi-product clock stage.

Separation of routes. Because undersold units are offered at regional prices that may lie below the package-implied price, a bidder that released a package unit could in principle exit and re-acquire the units it stranded at the lower price. Within the clock stage this is foreclosed by the no-conversion rule and separate eligibility: a released package unit does not become regional-level demand, and package eligibility cannot be used for regional-level bids. The eligibility rule above extends the same separation to the supplementary stage: a bidder bids in it only with regional eligibility acquired at the outset.

5.2 Assignment stage

The clock stage and supplementary stage allocate generic blocks. The assignment stage, which follows the design used in recent clock auctions, determines which specific blocks (A, B, …, X) each winner receives.

Because all S blocks in a region are similar, the assignment preferences of consequence are contiguity within a region, for a winner of more than one block, and consistency across regions, for a package winner or multi-region winner, for whom the same block letter everywhere simplifies network engineering. The stage therefore has two parts. In the *national part*, winners of package units bid for block assignments uniform across all regions. In the *regional part*, winners of regional units are offered assignments among the remaining blocks, contiguous ones for winners of more than one block; regions in which the same bidders won the same quantities may be combined so that a bidder can bid once for a consistent assignment across them.

In each part the system chooses the assignment maximizing the sum of submitted bids, breaking ties randomly, and charges assignment payments by the bidder-optimal core, nearest-Vickrey rule used in recent assignment stages. Assignment payments are additional to clock-stage payments, never negative, and zero for a bidder submitting no assignment bids. Blocks won in the supplementary stage are assigned with the regional winnings; the assignment stage does not distinguish blocks won through the package from blocks won regionally.

6. Properties of the Format

This section collects what can be proved about the format. Section 6.1 states the invariants that hold for every sequence of bids. Sections 6.2–6.5 analyze straightforward bidding in the LLG model: one package (Theorem 1), the prices it produces (Theorem 2), and several packages (Theorem 3). Section 6.6 places HPB in the same model and compares the two formats; Section 6.7 introduces budget constraints (Proposition 4); and Section 6.8 discusses the design choices on which the

results depend, including the incentive properties of the supplementary stage (Proposition 3). All proofs are in Appendix A.

6.1 Basic properties

The following properties hold for every sequence of bids under the specification of Sections 4 and 5; each follows directly from the structure of bid processing.

Proposition 1 (Basic properties). In the clock stage of the HCA:

(a) *Prices never fall.* For every region, $p_j^t \geq p_j^{t-1}$, and $p_0^t \geq p_0^{t-1}$, in every round.

(b) *The premium is non-negative and is absorbed.* $p_0^t \geq \sum_j p_j^t$ in every round. The premium ψ^t starts at zero, can increase only in a round in which the package ends with excess demand, and in every other round satisfies $\psi^t = \max\{\psi^{t-1} - \sum_j \left(p_j^t - p_j^{t-1}\right),\ 0\}$. In particular, if $D_0 \leq S_0$ at the end of every round, every package winner pays exactly the sum of the regional prices.

(c) *Bounded liability.* If a regional reduction at price point θ in region j is not applied, the posted price p_j^t is at most the regional price at θ. If a package reduction is not applied, the posted package price is the start-of-round price, $p_0^t = p_0^{t-1}$.

(d) *The reserve.* At the close, at most S_0 blocks in any region are won through packages. If regional-level demand in region j has ever reached $d \leq K$ units, it never falls below d thereafter; hence every reserved unit that is demanded at the regional level at its final posted price is won by a regional-level bidder.

(e) *Undersell.* A block is unsold at the close only if it is a reserved unit for which regional-level demand never reached the reserve, or capped aggregate demand for its product fell below supply through a package release applied while $D_0 \leq S_0$.

(f) *Pricing and termination.* All regional-level winners in a region pay the same price; every package winner pays $\sum_j p_j^T + \psi_T$; and the clock stage ends in finitely many rounds.

Two consequences deserve comment. Property (c) is the HCA's replacement for HPB's treatment of a displaced package bidder: a held bidder is never charged a price it did not accept, and a held package bidder is charged exactly the price at which it asked to leave. Property (b) says the premium is a transient, created by package competition and dissipated by regional competition: a package bidder that outlasts its package rivals and then faces rising regional prices pays the larger of the two, never the sum.

6.2 The local–local–global (LLG) model

To study the format under straightforward bidding we use a many-region, multi-unit version of the local–local–global model introduced by Krishna and Rosenthal (1996), in which bidders with single-item interests compete against a bidder whose value is for the package of all items. The model is the standard test environment for package auctions; Ausubel and Baranov (2020) use it to study core-selecting auctions under incomplete information. We refer to it as the LLG model.

The environment is as follows. There are J regions. Each region has $S = K + 1$ blocks, K of which are reserved, so that $S_0 = 1$: at most one package can be won. There are two kinds of bidders.

- A *local bidder* $\ell \in L_j$ has value $v_\ell > 0$ for one block in region j and no value for anything else. It bids only at the regional level and holds regional eligibility for one block in j.
- A *global bidder* $g \in G$ has value $V_g > 0$ for the package—one block in every region—and no value for any other bundle; a partial footprint is worthless to it. It bids only for the package and holds package eligibility for one unit.

Opening prices are zero, and G is non-empty. Write $v_j^{(i)}$ for the i-th highest value among the local bidders in region j, with $v_j^{(i)} = 0$ if $|L_j| < i$; write $V^{(1)} > V^{(2)}$ for the two highest global values, with $V^{(2)} = 0$ if $|G| = 1$, and g_1 for the global bidder with value $V^{(1)}$. Let

$$\Delta = \sum_{j \in R} v_j^{(K+1)}$$

denote the *displacement value* of the package: the sum over regions of the value of the local bidder that the package displaces.

A feasible allocation respects the reserve (at most one package is sold) and the bidders' unit demands. Since every value is positive, an efficient allocation either sells the package to g_1 and gives the K remaining blocks in each region to the K highest-value local bidders there, for total value $W_1 = V^{(1)} + \sum_j \sum_{i \le K} v_j^{(i)}$, or sells no package and gives the $K + 1$ blocks in each region to the $K + 1$ highest-value local bidders there, for total value $W_0 = \sum_j \sum_{i \le K+1} v_j^{(i)}$. Hence the package should be sold if and only if

$$V^{(1)} > \Delta.$$

We assume genericity throughout: values are distinct within each region and within G, and no global bidder's value equals the package price at any instant at which a local bidder exits. The last condition fails only on a set of value profiles of Lebesgue measure zero; it ensures that no global exit coincides with a local exit. (Local exits in different regions may coincide; they are processed independently and in either order, and nothing below depends on that order.)

To state an exact result we work with the continuous-clock idealization of the HCA, in which increments are vanishingly small and time is continuous. Prices start at zero. Let $A_j(t) \subseteq L_j$ be the set of active local bidders in region j, $A_0(t) \subseteq G$ the set of active global bidders, $D_j(t) = |A_j(t)|$ and $D_0(t) = |A_0(t)|$. Region j is *hot* at t if $\overline{AD}_j(t) = D_j(t) + \min\{D_0(t), 1\} > K + 1$; let $H(t)$ be the set of hot regions. Between events:

- p_j rises at rate one while $j \in H(t)$ and is constant otherwise (this is the regional clock);
- if $D_0(t) \geq 2$ and $H(t) = \varnothing$, p_0 rises at rate one (this is case (b) of Section 4.4); otherwise $p_0(t) = \max\{\Pi(t), \sum_j p_j(t)\}$, where $\Pi(t)$ is the package price at the end of the most recent own-increment interval, or zero if there has been none (this is case (a) of Section 4.4

together with the sum rule of Section 4.8: the package price is constant except that it rises with the regional sum whenever the two are equal).

A local bidder's *exit* at time t (a reduction of its demand to zero) is applied if and only if $j \in H(t^-)$; a global bidder's exit is applied if and only if $D_0(t^-) \geq 2$ or $H(t^-) \neq \varnothing$ (the exit rule). An exit that is not applied is held and re-tested at every subsequent event. The clock stage ends at the first time T at which $H(T) = \varnothing$ and $D_0(T) \leq 1$. The winners are the active bidders at T; local winners pay $p_j(T)$ and a global winner pays $p_0(T)$. Regions with $u_j = K + 1 - D_j(T) - D_0(T) > 0$ have undersell, and the parallel supplementary stage of Section 5.1 then runs a single-product clock in each such region, opening at $p_j(T)$, in which any local bidder not already a winner may demand one block; its price rises while demand exceeds u_j and it ends when it does not.

We use the term *straightforward bidding* in the sense of Milgrom (2000): a bidder demands, at every price, the quantity that is myopically optimal for it at that price. A local bidder ℓ demands one block in its region at every price $p_j \leq v_\ell$ and none above: it exits at the instant p_j reaches v_ℓ, and in the supplementary stage it demands a block at every price not exceeding v_ℓ. A global bidder g demands the package at every package price $p_0 \leq V_g$ and none above: it exits at the instant p_0 reaches V_g.

6.3 Straightforward bidding is efficient

Theorem 1. In the LLG model with generic values, suppose every bidder bids straightforwardly in the continuous-clock HCA with the parallel supplementary stage. Then:

(i) *No bidder is ever held.* Every exit is applied at the instant it is submitted, and the clock stage ends in finite time.

(ii) *If $V^{(1)} > \Delta$, the package is sold.* Global bidder g_1 wins the package at price

$$p_0^* = \max\{V^{(2)}, \Delta\},$$

and in each region the K highest-value local bidders win the remaining blocks at price $p_j^* = v_j^{(K+1)}$. No block is undersold except reserved units in regions with fewer than K local bidders.

(iii) *If $V^{(1)} < \Delta$, no package is sold,* and in each region the $K + 1$ highest-value local bidders win its blocks (all of them, if there are fewer than $K + 1$; the blocks no local bidder demands remain unsold). In a region that is still hot at the instant g_1 exits, they win in the clock stage, at price $v_j^{(K+2)}$ if the region has at least $K + 2$ local bidders; in a region that has already settled with K active local bidders, K of them win in the clock stage at $v_j^{(K+1)}$ and the $(K + 1)$-th wins the stranded block in the supplementary stage at the same price.

(iv) *Efficiency.* In either case the final allocation maximizes total value among allocations that respect the reserve.

(v) *Vickrey payments.* When the package is sold, every winner's payment equals its Vickrey–Clarke–Groves payment for the efficient allocation.

6.4 Core outcomes, Vickrey bounds and Walrasian prices

Theorem 1 determines the allocation and, when the package is sold, the payments. We now record what the payments satisfy in every case, assuming throughout this subsection that every region has at least $K + 2$ local bidders, so that every block is demanded and each order statistic below is an actual bidder's value.

The first benchmark is the coalitional game. For a set C of bidders, let $W(C)$ be the largest total value obtainable by allocating the blocks among the members of C alone, respecting the reserve and the bidders' demands; thus $W(N) = \max\{W_1, W_0\}$ for the set N of all bidders. An outcome with seller revenue R and bidder payoffs u_i (value less payment for a winner, zero for a loser) is in the *core* if it is efficient and no coalition of the seller and a set C of bidders could do better on its own:

$$R + \sum_{i \in C} u_i \;\geq\; W(C) \qquad \text{for every } C \subseteq N.$$

The second benchmark is Walrasian. Treat the blocks of region j and the package as $J + 1$ distinct goods—the *hierarchical economy*— with prices $p_1, \ldots, p_J$ and p_0. A local bidder in region j demands one block at $p_j < v_\ell$, none at $p_j > v_\ell$, either at equality; a global bidder demands the package at $p_0 < V_g$, none at $p_0 > V_g$, either at equality. The seller supplies $m \in \{0,1\}$ packages and $K + 1 - m$ blocks per region, choosing m to maximize revenue at the going prices: $m = 1$ if $p_0 > \sum_j p_j$, $m = 0$ if $p_0 < \sum_j p_j$, either if equal. A *Walrasian equilibrium* is a price vector (p, p_0) and an allocation in which every bidder receives a bundle it demands at those prices, the seller's supply is revenue-maximizing, and every good is fully allocated or has a zero price.

Theorem 2 (Prices). Under the hypotheses of Theorem 1, and with at least $K + 2$ local bidders in every region:

(a) *Core.* The outcome is in the core.

(b) *Vickrey bounds.* Every winner's payment is at least its Vickrey–Clarke–Groves payment for the efficient allocation, with equality for every winner when the package is sold. In particular, revenue is at least the Vickrey revenue.

(c) *Walrasian prices.* The final prices, together with the final allocation, form a Walrasian equilibrium of the hierarchical economy. When the package is sold they are the smallest Walrasian prices; when it is not, they need not be.

Two remarks are in order. First, (a) and (b) say that the HCA produces, by an ascending process, a core outcome whose payments dominate the Vickrey payments—the property that the core-selecting package auctions of Day and Milgrom (2008) are designed to deliver in one shot—and that it produces the Vickrey outcome itself whenever the package is sold. Second, (c) explains the shape of the result: in the hierarchical economy the package and the regional blocks are distinct goods, the seller chooses between them at the prices, and the HCA is a tâtonnement in that economy. The theorem needs $S_0 = 1$ because for one package the set of Walrasian prices is reached from below whatever the order in which regions settle; Section 6.5 shows why this fails with several packages, and what survives.

6.5 Several packages

Theorem 1 is stated for $S_0 = 1$, and the restriction is substantive. We now extend the model to $S_0 \geq 1$ packages. What survives: the HCA never sells too many packages, allocates efficiently given the number it sells, and is exact in both boundary cases; its only possible error is to sell too few, and the error has a simple diagnosis and a bound.

The environment generalizes that of Section 6.2. Each region has $S = K + S_0$ blocks, K of them reserved, so up to S_0 packages can be won. Local bidders are as before; a global bidder values one package at V_g and nothing else. Write $V^{(1)} > V^{(2)} > \cdots$ for the global values in decreasing order, with $V^{(m)} = 0$ when $|G| < m$, and assume every region has at least $S + 1$ local bidders. For $m = 1, \ldots, S_0$ define the *displacement value of the m-th package*,

$$\Delta_m \;=\; \sum_{j \in R} v_j^{(S-m+1)},$$

the sum over regions of the value of the local bidder displaced when an m-th package is sold, given that $m - 1$ are; $\Delta_1 = \sum_j v_j^{(S)}$, $\Delta_{S_0} = \sum_j v_j^{(K+1)}$. Selling m packages rather than $m - 1$ changes total value by $V^{(m)} - \Delta_m$, decreasing in m, so the efficient number of packages is

$$m^* \;=\; \max\{\, m \leq S_0 : V^{(m)} > \Delta_m \,\}$$

(with $m^* = 0$ if the set is empty); the efficient allocation gives the packages to the m^* highest-value global bidders and each region's remaining $S - m^*$ blocks to its $S - m^*$ highest-value local bidders. Genericity is as in Section 6.2, with $V^{(m)} \neq \Delta_m$ for all m.

The continuous clock is as in Section 6.2 with S_0 in place of 1: region j is hot if $D_j + \min\{D_0, S_0\} > S$; the package price rises on its own when $D_0 > S_0$ and no region is hot, and otherwise equals $\max\{\Pi, \sum_j p_j\}$; a global exit is applied if and only if $D_0 > S_0$ or some region is hot; the stage ends when no region is hot and $D_0 \leq S_0$. Straightforward bidding is as before. For the supplementary stage we take the parallel version of Section 5.1, but with the clock in each undersold region opening at the clock-stage *opening* price (zero) rather than at the final posted price; the reason is explained after the theorem.

Theorem 3 (Several packages). In the environment above, under straightforward bidding in the continuous-clock HCA, let $\hat{m}$ be the number of packages sold. Then:

(i) no exit is ever held, and the clock stage ends in finite time;

(ii) *one-sided error:* $\hat{m} \leq m^*$, and the packages are won by the $\hat{m}$ highest-value global bidders;

(iii) *exactness at the boundaries:* if $m^* = S_0$ then $\hat{m} = S_0$, each global winner pays $\max\{V^{(S_0+1)}, \Delta_{S_0}\}$, the K highest-value local bidders in each region win its remaining blocks at $v_j^{(K+1)}$, and every payment is the Vickrey payment; if $m^* = 0$ then $\hat{m} = 0$;

(iv) *conditional efficiency and the size of the error:* in every case the $S - \hat{m}$ highest-value local bidders in each region win its remaining blocks, in the clock stage or the supplementary

stage, so the allocation is efficient given $\hat{m}$; and the welfare loss relative to the efficient allocation is

$$\sum_{m=\hat{m}+1}^{m^*} \left(V^{(m)} - \Delta_m\right) \;\le\; (m^* - \hat{m}) \sum_{j \in R} \left(v_j^{(K+1)} - v_j^{(S-m^*+1)}\right).$$

The reason the error is one-sided is instructive. An ascending clock reveals each region's local values from the bottom: $v_j^{(K+2)}$ is passed before $v_j^{(K+1)}$. The decisions about the number of packages need them in the opposite order: whether to keep the S_0-th package requires $v_j^{(K+1)}$, which the clock reaches last; whether to keep the $(S_0 - 1)$-th requires $v_j^{(K+2)}$, which the clock has already passed and, prices being monotone, cannot revisit. A region that settles while S_0 packages are still demanded stops at $v_j^{(K+1)}$; when a package later departs, the region's contribution to the package price stays there although the relevant displacement value has fallen to $v_j^{(K+2)}$, and a remaining global bidder may be driven out by a regional sum that overstates the value of what it displaces. For example, with $J = 2$, $K = 1$ and $S = 3$ ($S_0 = 2$), local values (10,9,8.5) in region 1 and (10,3,2) in region 2, and global values 11.2 and 11, it is efficient to sell exactly one package (since $11.2 > 8.5 + 2$ but $11 < 9 + 3$); yet region 2 settles at 3 while both global bidders are present, the second global bidder leaves when the package price reaches 11, and the remaining global bidder then faces a regional sum that rises to $8.5 + 3 = 11.5 > 11.2$ and leaves as well, although its value exceeds the displacement value 10.5 of the single package it would win. Under-selling is the only possible error because stale prices are always too high, never too low.

The failure is one of the process, not of the existence of equilibrium prices: a Walrasian equilibrium exists for every S_0. Extend the hierarchical economy of Section 6.4 by letting the seller supply $m \in \{0, \dots, S_0\}$ packages and $S - m$ blocks per region, choosing m to maximize revenue. A Walrasian equilibrium with m packages requires $p_j \in \left[v_j^{(S-m+1)}, v_j^{(S-m)}\right]$ in every region and, for the global bidders and the seller, $\max\{V^{(m+1)}, \Delta_m\} \le \sum_j p_j \le \min\{V^{(m)}, \Delta_{m+1}\}$ (with the obvious conventions at $m = 0$ and $m = S_0$). This window is non-empty exactly when $m = m^*$, so the economy always has a Walrasian equilibrium, with m^* packages. For $S_0 = 1$ the only constraints a rising regional sum can violate are lower bounds ($\sum_j p_j \ge \Delta_1$ when $m^* = 1$, $\sum_j p_j \ge V^{(1)}$ when $m^* = 0$), and a rising sum satisfies them eventually, which is why Theorem 1 is exact; for an interior m^* the window has an upper bound, $\min\{V^{(m^*)}, \Delta_{m^*+1}\}$, and a region that settled under a higher package count contributes more than its current displacement value, so the regional sum can exceed $V^{(m^*)}$ although Δ_{m^*} does not.

The bound in (iv) also says when the error is small. It is the spread, summed over regions, between the local bidder displaced by the S_0-th package and the one displaced by the m^*-th. If a region's local values are drawn independently from a distribution with density bounded away from zero and the number of local bidders grows, consecutive order statistics are within $O_p\left(1/n_j\right)$ of one another, so the bound, and with it the welfare loss, vanishes as regional markets thicken. An exact remedy exists but we do not adopt it: the auctioneer knows the price at which each local bidder left, and a variant in which a region's contribution to the package price reverts, when a package departs

and strands a block there, to the exit price of the highest-value inactive local bidder (kept monotone by using it only as a floor) restores $\hat{m} = m^*$. That variant prices the same region's blocks differently to different winners and is, in substance, revival at the exit price; the trade-off between exactness and the clean price structure is real, and we prefer the latter.

A final remark concerns the supplementary stage. If two or more packages depart after a region has settled, it strands two or more blocks, but only the local bidder whose exit set $p_j(T)$ is willing to buy at $p_j(T)$. Conditional efficiency in (iv) therefore needs the supplementary clock to open below the final posted price—at the opening price, as stated—so that the second stranded block can find its buyer. With $S_0 = 1$ at most one block is stranded per region and the two versions coincide, apart from the incentive question of Section 6.8.

6.6 HPB in the same model

The claim of this paper is that the HCA improves upon HPB. This section places HPB in the environment of Theorem 1 and compares the two formats under the same behavioral assumption. The comparison is instructive twice over: it shows where HPB's inefficiencies come from, and that they are not where one might first look.

HPB must first be translated into the LLG environment. It is defined over specific licenses, so region j's $K + 1$ blocks become $K + 1$ licenses: one *open license* a_j, belonging to the national package $A = \{a_1, \dots, a_J\}$, and K *reserved licenses* $b_j^1, \dots, b_j^K$, which can be bid for only individually and which the package cannot win. Local bidders regard their region's $K + 1$ licenses as perfect substitutes; global bidders bid only for A. The rules are those of the 2008 auction (FCC, 2007; Goeree and Holt, 2010), reduced to this setting:

- Bids are standing offers. A bid remains live until it is outbid, and cannot be withdrawn (or, where withdrawal is permitted, only at a withdrawal payment).
- After each round, the package bid B is compared with the sum $\sum_j h\left(a_j\right)$ of the high bids on its constituents; the higher is provisionally winning (the licenses win ties). Free licenses are always provisionally won by their high bidders. A bidder that is provisionally winning on a license does not bid on it again.
- The minimum acceptable bid on a reserved license, or on a open license while the licenses are provisionally winning, is its high bid plus the item increment δ. While the package is provisionally winning, the minimum acceptable bid on a_j is its *current price estimate* plus δ, where the estimate is $h\left(a_j\right) + s_j\,\sigma$, $\sigma = B - \sum_k h\left(a_k\right)$ is the shortfall and s_j is region j's share of the package's bidding units. The minimum acceptable bid on the package is $\sum_j h\left(a_j\right) + \delta_0$ while the licenses are provisionally winning, where δ_0 is the package increment.
- A bidder with no provisionally winning bid and no new bid loses activity and, once its waivers are exhausted, its eligibility. The auction ends after the first round with no new bid.

Straightforward bidding is the natural counterpart of Section 6.2: in every round, a local bidder that is not provisionally winning any license bids the minimum acceptable bid on the cheapest license in its region (the one with the lowest minimum acceptable bid), provided that bid does not exceed its

value, and otherwise bids nothing; a global bidder bids the minimum acceptable bid on the package whenever the package is not provisionally winning and that bid does not exceed its value. Ties between equally cheap licenses are broken by a fixed rule. Under this profile the package bid never exceeds the item sum by more than δ_0, so the shortfall σ is at most one package increment and the only frictions are the increments themselves. That is the point: the comparison with the HCA is then between the frictions of the two processes, not between two behavioral assumptions.

Proposition 2 (HPB in the LLG model). Fix J, K, the shares s_j and increments $\delta, \delta_0 > 0$.

(a) *HCA.* Under straightforward bidding in the HCA with a positive increment, the allocation does not depend on the order in which bids are processed or on tie-breaking, and it differs from the efficient allocation only if $V^{(1)}$ lies within one package increment of Δ or of $V^{(2)}$. A regional bidder's exit is always applied at the price point equal to its value: regional bidders face no increment friction at all.

(b) *HPB: path dependence and exposure.* Under straightforward bidding in HPB, the allocation depends on the order of bids and on the tie-breaking rule, and for every $\delta, \delta_0 > 0$ there is an open set of value profiles with $\Delta > V^{(1)}$ on which, for some tie-breaking rule, a local bidder wins two licenses—one of them through revival of a bid it placed before it won the other—pays more than its value in total, and a local bidder whose value exceeds the final price in its region wins nothing. The welfare loss equals the excluded bidder's value.

(c) *HPB: apportionment.* While the package is provisionally winning with shortfall σ, a local bidder in region j can keep the package contested only by bidding at least $h(a_j) + s_j\sigma + \delta$ on the open license; a region whose marginal local bidder cannot afford its apportioned share stops contributing to the item sum at a level up to $s_j\delta_0 + \delta$ below its displacement value $v_j^{(K+1)}$, and must then be carried by the other regions. Consequently HPB can sell the package on value profiles with $\Delta > V^{(1)}$, the set of such profiles having width of the order of $\max_j s_j\delta_0$ in each region's displacement value.

A worked example. The following example establishes part (b) of Proposition 2; the set of profiles is open because every inequality used is strict. Let $J = 2$, $K = 1$, $s_1 = s_2 = \frac{1}{2}$, $\delta = 1$, $\delta_0 = 2$, and opening bids of 1 on every license and 2 on the package. Region 1 has local bidders ℓ_1 and ℓ_2 with values 10 and 7.9; region 2 has m_1 and m_2 with values 30 and 28; the single global bidder has value $V = 24$. Since $\Delta = 7.9 + 28 = 35.9 > 24$, the efficient allocation sells no package and gives each local bidder one license, for a total value of 75.9. Ties are broken as follows: ℓ_1 and m_1 prefer the open license, ℓ_2 and m_2 the reserved license. In round 1, ℓ_1 bids 1 on a_1, ℓ_2 on b_1, m_1 on a_2, m_2 on b_2, and the global bidder bids 2 on the package; the licenses win the tie. The auction then runs in four-round cycles. In the first round of a cycle the global bidder, displaced, bids the item sum plus 2 and the package becomes provisionally winning, so the holders of a_1 and a_2 lose that status and the minimum acceptable bid on each open license rises by 2 (its own increment plus half the shortfall) while that on each reserved license rises by 1. In the second round ℓ_1 and m_1, whose cheapest option is now the reserved license, take b_1 and b_2 at one increment above their holders. In the third round the displaced ℓ_2 and m_2 face equal minimum bids on the two licenses and, by their tie-break, retake b_1 and b_2. In the fourth round ℓ_1 and m_1, for whom the open license is now cheaper, return

to a_1 and a_2 at two increments above their own dormant bids; the item sum exceeds the package bid and the licenses are provisionally winning again. Each cycle raises every high bid by 2; after three cycles all four high bids are 7 and the package bid is 12. In the fourth cycle the global bidder bids 16, ℓ_1 takes b_1 at 8, and ℓ_2 is now unable to bid: the reserved license requires 9 and the open license 9 (the dormant bid of 7, plus 1 for its share of the shortfall, plus δ), both above 7.9. Region 1 is thereafter silent: ℓ_1 is provisionally winning b_1 at 8 and holds the high bid of 7 on a_1, and ℓ_2, having no provisionally winning bid and placing none, loses its eligibility. In region 2, m_1 and m_2 continue the cycle, and each time the item sum overtakes the package bid ℓ_1's dormant bid on a_1 is revived. The global bidder's bids run 16,18,20,22,24; when the item sum next reaches 24 (with $h(a_1) = 7$, $h(a_2) = 17$), the next package bid would be $26 > 24$, and the auction ends. The final allocation gives ℓ_1 both licenses of region 1, at $7 + 8 = 15$ for a value of 10; ℓ_2 nothing, although the last price at which a_1 changed hands was $7 < 7.9$ and its minimum acceptable bid at the close, 8, is also above 7.9; and m_1 and m_2 one license each at 17. Total value is 68, a loss of 7.9. Under the HCA with the same values, Theorem 1 gives each local bidder one license: region 1 settles at 7.9 while the package is present, the global bidder leaves when the regional sum reaches 24 (region 2's price is then 16.1), region 2 settles at that price, and ℓ_2 recovers the stranded block at 7.9 in the supplementary stage. $

The comparison yields three observations.

First, the source of HPB's inefficiency in this environment is not the one usually named. The apportionment rule is an increment effect: as δ and δ_0 vanish, a region's frozen contribution approaches its displacement value and (c) disappears. So does the exposure of (b): the dormant bid that revives is always within $s_j\delta_0 + \delta$ of the value of the bidder it excludes. In the limit, straightforward bidding in HPB also drives the item sum to Δ, displaces the package exactly when $\Delta > V^{(1)}$, and hands each open license to a bidder who values it at its bid. We record this as a conjecture rather than a theorem, HPB's path dependence making a general limit statement awkward; but we have found no first-order inefficiency of HPB under straightforward bidding in this model, and we believe there is none. The threshold problem proper—each regional bidder's incentive to let the others cover the shortfall—is a property of equilibrium behavior, not of straightforward bidding, and lies outside this comparison.

Second, the increment effects are nonetheless the right comparison, because their scales differ. In HPB the increment a regional bidder must absorb to keep the package contested is its region's share of a *package* increment, $s_j\delta_0$, apportioned by bidding units rather than by value: in an auction with a multi-billion-dollar package and percentage increments, a region with one percent of the bidding units and a ten percent package increment is asked for millions of dollars between one round and the next, whatever its own licenses are worth; regions whose value share is below their bidding-unit share freeze first, and the rest must carry them. In the HCA a regional bidder never meets a package increment: its region's contribution to the package price is its own regional price, it exits at its value exactly, and the only increment in the system is borne by the global bidder, at the scale of the package it bids for. The worked example is small only because the increments are; the mechanism scales with the package.

Third, independently of increments, the two formats differ in what determines the allocation. Under the HCA it is a function of the values (Proposition 2(a) and Theorem 1); under HPB, of the values *and* the order in which bids happened to be placed, and on some orders a regional bidder is exposed. This is the precise form of the liability problem of Section 1: because the open and reserved licenses are substitutes for a local bidder, and bids are standing offers, a bidder that bids on the open license and is then displaced by the package must either stop bidding—and risk exclusion—or bid on a substitute and risk winning both when its dormant bid revives. HPB removes exposure for the package bidder and relocates it to the regional bidders. The HCA removes it for both: a regional bidder's liability is its current demand at the current price (Proposition 1(c)), and the reserved units, not revival, keep regional demand relevant.

It is natural to ask where the 2008 auction itself stands in this environment. The C block of the 2008 U.S. 700 MHz auction consisted of twelve regional licenses, one per region, with a hierarchy of pre-defined packages above them, the largest covering the fifty states (FCC, 2007). In the notation of Section 6.2 this is the case $K = 0$: one block per region, no reserve. Theorem 1's argument does not use $K \geq 1$ and applies (the HCA then has no reserved units; the package exit rule and the supplementary stage do the work). With $K = 0$ the within-region substitutes driving Proposition 2(b) are absent from the hierarchy, but in 2008 they were present outside it—the A and B block licenses covering the same areas, sold in the same auction under ordinary SMR rules—so the same mechanism operates across blocks: a bidder with a dormant bid on a C block license that then wins A or B block licenses in the same area is exposed to the former's revival. Proposition 2(c) applies unchanged, and with percentage increments on a package bid in the billions its scale is the one just described.

6.7 Budget constraints

The introduction listed, among the consequences of bid revival, the difficulty it creates for a budget-constrained bidder: a live bid encumbers funds whether or not it is winning. This section makes the point formally. A unit-demand bidder with value v and budget w simply behaves as one with value $\min\{v, w\}$, so budgets have no bite in the model of Theorem 1; the natural environment has a constrained bidder demanding blocks in several regions, so that its liabilities can accumulate.

Modify the LLG model as follows, keeping $S = 2$ and $K = 1$. Each region j has one local bidder ℓ_j with additive values: each of the two blocks is worth θ_j to it, so it demands two blocks at any price below θ_j and none above. There is one global bidder with value V for a single package, and one *cross-region bidder* β, with value v_β for one block in each region of a set $R_\beta \subsetneq R$, $|R_\beta| = 2$, additive across regions, and a hard budget w: under no circumstances can β pay more than w in total. All other budgets exceed any conceivable payment; θ_j and V are private information; genericity holds as before. A strategy of β's is *budget-feasible* if, for every behavior of the others, every processing order and every tie-break, β's payment obligations never exceed w.

Two observations organize the analysis. First, the displacement value of the package is $\Delta = \sum_{j \in R} \theta_j$ in every state: with or without β, the block the package displaces in region j is one ℓ_j would otherwise win, since ℓ_j demands both. So Δ does not depend on v_β or w, and the package decision is unaffected by anything β does. Second, the two formats account for β's liabilities differently. In the

HCA, β's liability at any moment is its current processed demand at current prices—observable, and controlled by β directly. Under HPB it is the sum of β's *live high bids*, because a package displacement revives every dormant bid beneath it at once; a dormant bid encumbers the budget while purchasing nothing unless revived.

Proposition 4 (Budget constraints). In the environment above:

(a) *HCA.* The truncation strategy—demand one unit at the regional level in each region of R_β while $p_j \leq v_\beta$ and $\sum_{k \in R_\beta} p_k \leq w$; if the budget constraint is about to bind, reduce to zero in a region of R_β whose price is rising, choosing one with the smallest $v_\beta - p_j$—is budget-feasible. If β follows it and all other bidders bid straightforwardly, then for every realization of (θ, V), every processing order and every tie-break, the clock stage (with the supplementary stage) yields the efficient allocation subject to β's constraint: the package is sold if and only if $V > \Delta$; each region's blocks are allocated efficiently given the package decision and the regions β retains; and β pays at most w. The strategy requires no knowledge of (θ, V), and the package decision and the payoffs of all other bidders are the same as if β were unconstrained whenever $w \geq \sum_{k \in R_\beta} \theta_k$.

(b) *HPB.* In the HPB counterpart of Section 6.6: (i) straightforward bidding is not budget-feasible whenever $w < 2 \sum_{k \in R_\beta} \min\{v_\beta, \theta_k\}$ approximately (up to increments and opening bids), even when w comfortably exceeds β's efficient bill $\sum_{k \in R_\beta} \theta_k$: on orders of play of the kind constructed in Proposition 2(b), β comes to hold, simultaneously, a live high bid on both the open license and the reserved license of each region in R_β, so that its worst-case liability is roughly twice its intended purchase in every region at once. (ii) Under the SMR activity rule, there is an open set of realizations with $V < \Delta$, and orders of play consistent with budget-feasible bidding that follows straightforward bidding whenever it is feasible, on which β wins nothing although w exceeds its efficient bill and $v_\beta > \theta_j$ for every $j \in R_\beta$; the welfare loss, up to $\sum_{k \in R_\beta} (v_\beta - \theta_k)$, is bounded away from zero as the increments vanish.

Four remarks delimit the result. First, the activity rule is essential to (b)(ii), and deliberately so: without it, a patient bidder could wait for rivals to outbid its stale positions and re-enter as its encumbered funds are released, and the loss would shrink to the increment scale of Proposition 2. The toxic combination is standing bids, worst-case liability and activity requirements together: frozen capital costs eligibility, and lost eligibility makes the freeze permanent. In the HCA the same activity rule poses no difficulty, because β's activity *is* its live demand, which the truncation strategy keeps within budget throughout. Second, no welfare ranking of the formats is claimed. HPB offers β budget-feasible escapes—notably, bidding only for the reserved licenses, which carries no revival exposure and caps ℓ_j's willingness to chase at θ_j—but every escape abandons straightforward bidding for an ex ante, region-by-region choice of license whose safety depends on the package outcome. The differentiator is thus not that HPB must destroy welfare against sophisticated play, but that it converts a budget into a strategy problem, while the HCA converts it into an exit rule. Third, straightforward bidding by the additive local bidders is itself not an equilibrium of a uniform-price clock—multi-unit demand brings the familiar demand-reduction incentive (Ausubel,

Cramton, Pycia, Rostek and Weretka, 2014)—so, as with Proposition 2, the comparison is between the formats under a common behavioral benchmark, not between equilibria; Section 9.2 keeps both questions open. Fourth, the model assumes β cannot default. Under limited liability the mechanism inverts: revival generates surprise obligations rather than exclusion, a problem for the seller rather than for β. Budget constraints of both kinds were central to bidding in the 700 MHz auction itself (Bulow, Levin and Milgrom, 2009), the auction in which HPB was fielded; on auction design with financially constrained bidders generally, see Che and Gale (1998).

6.8 Discussion

The final package price is a competitive price in a precise sense. The winning global bidder pays the larger of two numbers—the value of its strongest package rival and the local bidders' aggregate displacement value—each the price that a different group of competitors would have to be paid to stand aside. The premium $\psi^* = \max\{V^{(2)} - \Delta, 0\}$ is positive exactly when global rather than local competition determines the package price; and the local bidders pay only their own region's displacement value, never a share of the global contest. This is the insulation property of Section 7.1, in equilibrium terms.

The results also show why the premium must be absorbed rather than carried. They depend on the rule of Section 4.4 under which the package clock price is the *larger* of the package's own price and the sum of the regional clock prices, so that regional price increases are absorbed into an existing premium. Consider the alternative rule under which the package price rises by the sum of the regional increments, carrying the premium unchanged. Take $J = 1$, $K = 1$ ($S = 2$), local bidders with values 10, 8 and 6, and global bidders with values 9 and 7. The efficient allocation sells the package, since $V^{(1)} = 9 > \Delta = 8$. Under the carrying rule, suppose the global contest ends at $p_0 = 7$ when the regional price stands at 3, so that the premium is 4; the package price is then $p_1 + 4$, reaches 9 when $p_1 = 5$, and g_1 exits—although the regional price would have had to reach 8 to displace it on the merits. Under the absorbing rule the package price is $\max\{7, p_1\}$, reaches 8 only when $p_1 = 8$, and g_1 wins at $8 = \max\{V^{(2)}, \Delta\}$.

The supplementary stage is best understood as revival at current prices, and the service has an incentive cost. In case (iii) of Theorem 1, the local bidder with value $v_j^{(K+1)}$ in a settled region left the clock stage while the package was present and one fewer block was available; when the package later departs, the block it would have won is stranded, and the supplementary stage returns it to that bidder at the price at which it left. This is precisely the service bid revival performs under HPB—but performed at the current price, for a bidder demanding the block now, rather than by reinstating a stale bid. Without the supplementary stage, the clock stage alone is efficient in case (ii) and falls short in case (iii) by exactly $\sum v_j^{(K+1)}$ over the regions settled before g_1 left. The service has a cost, which the next result makes precise. Say a strategy profile is an *ex post equilibrium* if, for every profile of values, no bidder can gain by deviating given that the others follow it.

Proposition 3 (Incentives in the clock stage and the supplementary stage). In the LLG model with generic values:

(a) In the continuous-clock HCA *without* a supplementary stage, straightforward bidding is an ex post equilibrium.

(b) With the parallel supplementary stage opening at the final posted prices, it is not: in every region that settles before g_1 exits, the local bidder with value $v_j^{(K+1)}$ can gain by exiting below its value, and the gain is positive on an open set of value profiles.

The deviation in (b) is the mild form of revival's incentive problem: the bidder whose exit sets a region's price gets the block back at that price, so it has reason to set the price low. The shading is bounded—ℓ loses the block altogether if the package then survives—and its direction is to understate displacement values, that is, toward selling the package when it should not be sold. Opening the supplementary clock one increment above the final posted price, or at the opening price as in Section 6.5, removes the gain in (b) but recreates the undersell the stage exists to cure (the stranded block is no longer worth $v_j^{(K+1)}$ to anyone at a price above it). The supplementary stage thus trades efficiency against incentives, and the trade-off deserves a proposition of its own in future work. Proposition 3(a) also explains why Theorem 1 needs no equilibrium hypothesis: in the clock stage proper, straightforward bidding is what a bidder does when it expects others to do the same.

With a positive increment the discrete HCA differs from the continuous idealization only in the granularity of the global bidder's exit decision. A package reduction is submitted at the start-of-round price, so a global bidder exits at the first round whose package clock price exceeds its value; its exit is applied, by the second remark of Section 4.7, provided it is the round's first package reduction—which, in the LLG model with $D_0 \leq 1$ after the global contest, it always is. Local bidders exit at price points equal to their values and face no granularity at all. The continuous-model allocation therefore obtains whenever no global value lies within one package increment of Δ or of another global value (Proposition 2(a)).

Losing-side prices, by contrast, are path-dependent. In case (iii) the allocation is determined by the values alone, but the regional prices are not: a region settled before g_1 left posts $v_j^{(K+1)}$, while a still-contested one posts $v_j^{(K+2)}$ or the price it had reached at t_1—and which regions had settled depends on the speed of the clocks. This familiar feature of simultaneous ascending auctions with monotone prices has no bearing on the allocation; it is why the case (iii) prices in Theorem 2(c) are Walrasian but not the smallest Walrasian prices.

Finally, several relaxations of the model appear to change nothing essential, and we omit them for brevity: local bidders with multi-unit demand and decreasing marginal values (read "marginal unit" for "local bidder" throughout, and define Δ with the $(K+1)$-th highest marginal value in each region); region-specific reserves K_j (Section 8.1); and bidders holding both a package demand and additively separable regional demands, which the separate-eligibility rule keeps apart. An inductive version for a Rothkopf–Pekeč–Harstad tree with the absorbing rule at every level (Section 8.2) is plausible but unproven, and the stale-price problem of Section 6.5 reappears across levels. What we would not conjecture without further work is anything about route choice—a bidder deciding at the

outset between the package and regional bidding—or about equilibrium with bidders of mixed type, which is where the premium's role as the price of exposure protection lives.

7. Design Rationale

The HCA is built from the elements of the standard clock format—generic products, start-of-round and clock prices, intra-round bidding, sequential processing under a no-excess-supply rule, eligibility-based activity, a separate assignment stage—and from the set-aside treatment in that format, transposed from bidder classes to bid types. Three choices go beyond that inheritance and deserve justification in the light of Section 6: a single regional product per region together with a package that has its own price; a reserve that restricts bids rather than bidders; and the package exit rule. We take them in turn.

7.1 Why one regional product, and a package with its own price

Each region has a single regional product because, with only a national package above it, there are exactly two ways to bid for a block in a region: at the regional level, which can win any block including the reserved ones, and through the package, which can win only the non-reserved blocks. The regional product carries the price of the first route; the package carries the price of the second. A separate "open" regional product would duplicate the package's price region by region—and would require a rule apportioning the premium among regions, precisely the apportionment the design is meant to avoid. Separate reserved and unreserved products per area are appropriate where reserved blocks are bid for area by area; they are not needed when the only unreserved route is a single nationwide bundle with a price of its own.

The two-case rule for the package clock price (Section 4.4) keeps the package price from running ahead of the regional prices. While any region has regional excess demand, the package's clock range is the sum of the regional clock ranges, so a package bid at a price point faces exactly the sum of the regional prices at that point, and the package's posted price follows the regional posted prices. Only when no region has regional excess demand is the package permitted its own percentage increment. The resulting premium is confined to competition among package bidders and never charged to a regional bidder; and because the package clock price is the *larger* of the package's own start-of-round price and the sum of the regional clock prices, later regional price increases are absorbed into the premium rather than added on . Section 6.8 shows that the alternative—carrying the premium and adding regional increases to it—can displace a package bidder whose value exceeds both its package rivals' values and the regional bidders' aggregate willingness to pay.

We considered, and rejected, a single price per region with package demand simply counted in the regional aggregate. That design would make the price of the reserved units hostage to package competition: package bidders contesting the available packages would drive up every region's price, a regional bidder demanding a reserved unit would be priced out of a unit no package bidder could win, and the unit would go unsold. The capped aggregate of Section 4.3 and the package's own price prevent this.

7.2 Why the reserve restricts bids rather than bidders

In the set-aside design, the set-aside is available only to a designated class of bidders. The HCA's reserve is available to any bidder willing to bid at the regional level and bear exposure. This choice has three consequences.

First, it neutralizes the *structural* advantage of package bids—the ability to sweep a region with an all-or-nothing bid—without favoring any class of firm. A national bidder that prefers to assemble its footprint region by region competes for the reserved units on equal terms with everyone else. For the same reason, the reserve does not mandate an inefficient allocation when the national use is more valuable: a bidder with national interests can win the reserved units region by region, bearing exposure, while winning the rest through the package.

Second, it makes the premium a price any bidder can pay or forgo—but the choice is made when demand is placed, not traded upon afterwards. Package demand and regional-level demand are never converted into one another (Section 4.2): a bidder holding the package cannot descend into regional-level demand in the regions it prefers while shedding the rest, and a bidder that has assembled a footprint cannot fold it into a package. The two routes are separate products with separate prices and separate eligibility, and a bidder wishing to keep both open must acquire eligibility for both at the outset. This keeps the package a genuine commitment, as a package bid is under HPB, and prevents the premium from being arbitraged away by bidders moving into and out of the package as its price moves relative to the regional prices. The premium is disciplined instead by competition among package bidders and by their *ex ante* choice of route: a bidder that regards exposure protection as worth less than the premium it anticipates will bid regionally from the outset. Section 5.1 extends the separation of routes to the supplementary stage as an option.

Third, it removes the costless waiting position that HPB affords regional bidders. Under HPB, an item bid beneath a provisionally winning package is dormant: it costs its owner nothing unless revived, so a regional bidder can sit on it and wait for the others to push the item sum past the package—the free-rider problem of Section 1. Under the HCA, regional-level demand is live demand at the current price. Because a regional reduction may be held by the no-excess-supply rule, and because at least K units in every region must go to regional-level demand (Proposition 1(c)–(d)), a regional bidder holding demand it does not want—to keep eligibility, or in the hope that others will displace the package—risks winning at the posted price. The risk is bounded, since a held bidder never pays more than the price at which it asked to reduce; but it is exactly the risk a bidder bidding its values accepts, so there is no position from which a regional bidder can wait for free. In the LLG model this is enough: straightforward bidding is an ex post equilibrium of the clock stage (Proposition 3(a)), with a residual, bounded incentive to shade in the supplementary stage (Proposition 3(b)). Whether the comparison with HPB survives in equilibrium—whether the HCA's incentives are better than HPB's once item bidders under HPB behave strategically—is open (Section 9.2).

7.3 Why the package exit rule

The no-excess-supply rule for regional products is what allows a standard clock auction to guarantee that a unit for which demand has reached supply will not go unsold, and it prevents a

bidder from pushing up a rival's price and then leaving the rival's unit to be sold at a lower price. For the package, three ways of applying the rule suggest themselves.

Strict application would release a package unit only if no region were left with regional excess supply. This deadlocks whenever a package bidder is the sole demander of a non-reserved unit in any region; and because the package price continues to rise while the bidder is held, its liability is unbounded.

Partial application would release the unit in regions with excess demand and convert the bidder's demand elsewhere into regional-level demand held at the posted price. This bounds liability but leaves a bidder that wanted to exit holding exactly the units no one else wants: a package bidder facing this rule would be no better off than one holding J regional products under the standard rules, and package bidding would protect no one. It also violates the principle that package demand is never converted into regional-level demand.

Unconditional release would apply every package reduction in full. This protects the package bidder completely, but lets it walk away from a package whose price is not rising—a package it is, at that moment, winning at a price it previously accepted—and strand units in the process.

The package exit rule takes the middle course: release in full, but only while the package price is being driven up, whether by package excess demand or by regional excess demand somewhere. A package bidder is held only when it is winning the package at exactly the start-of-round price at which it asked to leave (Proposition 1(c)), and can leave only when it is being outbid. This is the clock-auction analogue of the condition under which a package bidder is released from its commitment under HPB—its package bid ceases to be the PWB—translated from the language of standing bids into the language of demand at current prices.

The cost of the rule is undersell. Two observations put this in proportion. The reserved units in a region are never stranded by a package exit, because they were never the package's to hold. And a stranded unit is a unit that no bidder other than the departing package bidder demanded at its posted price—under HPB it would revert to the seller's bid in exactly the same circumstances—and the supplementary stage offers it once more before the assignment stage.

8. Variants and Extensions

8.1 Region-specific supply and reserves

The specification takes S and K common to all regions. If they vary, with S_j and K_j in region j, the package supply is $S_0 = \min_j(S_j - K_j)$, the regional excess-demand test becomes $D_j + \min\{D_0, S_j - K_j\} > S_j$, and the per-bidder limit becomes $x_{i,0} + x_{i,j} \leq S_j$. A larger reserve strengthens regional bidders at the cost of fewer winnable packages; note that with $S = 2$ and $K = 1$ at most one national package can be won, whereas HPB with two block-specific national packages would permit two.

8.2 An intermediate tier of regional packages (sketch)

The base-case clearing rules are worked out in full; those for an intermediate tier are not, and what follows is a sketch of the natural generalization, not a specification.

Let $\{G_1, \dots, G_m\}$ partition R, and let each G_j be a package product comprising one block in each of its regions. Package-covered demand in region j is $\widetilde{D}_j = D_0 + D_{G(j)}$, where $G(j)$ is the group containing j. The reserve requires $\widetilde{D}_j \le S_0$ in every region, so the package constraint is region-specific rather than global: a regional package has excess demand when $\widetilde{D}_j > S_0$ for some $j \in G_j$, and the national package when $\widetilde{D}_j > S_0$ for some j at all. The regional excess-demand test becomes $D_j + \min\{\widetilde{D}_j, S_0\} > S$, and the exit rule for a package releases a unit when the package's price is rising—when some covered region has $\widetilde{D}_j > S_0$ or regional excess demand.

The open questions concern prices. Each package would carry its own price, normally the sum of its constituents' prices, the national price normally the sum of the regional-package prices. When only package competition is present in a group—say two regional bidders contesting G_j with no regional excess demand there—the price of G_j must rise on its own, and the national price with it (since the national package contains G_j); when national bidders alone compete, the national premium must rise, but whether and how it should flow down to the regional packages is not obvious. The ordering of price points across three ranges, and the treatment of premia arising at different tiers, require further work.

8.3 Bidder-class eligibility for the reserve

If policy calls for it, the reserved units may additionally be restricted to a designated class of bidders. The only change is a validation check at submission limiting ineligible bidders' regional-level demand to the non-reserved units, which brings the design back to the set-aside treatment of bidder classes.

8.4 Package clock price summed over regions with excess demand only

Under the base rule the package's case-(a) clock range is the sum of *all* regional increments, although only regions with regional excess demand post above their start-of-round prices. When the package itself has excess demand, it therefore posts at a clock price including the increments of regions whose own prices did not move, the difference becoming premium to be absorbed later (Proposition 1(b)). A variant sets the case-(a) package clock price to the previous posted price plus the increments of the regions with regional excess demand at the end of the previous round, so the package price rises only at the pace of regional competition even when the package has excess demand. The cost: package excess demand coinciding with regional excess demand in only a few regions is resolved only at the pace of those regions' increments.

8.5 Threshold variants of the exit rule

The exit rule releases a package when a single region has regional excess demand. A stricter variant requires regional excess demand in regions whose bidding units sum to at least a fraction ϕ of b_0,

or a package clock price exceeding its start-of-round price by a stated percentage. Such variants make package demands more nearly the firm commitments that package bids are under HPB, at the cost of weakening Proposition 1(c). The base design sets $\phi = 0$, under which a held package bidder pays exactly the start-of-round price.

8.6 Supplementary-stage safeguards

If the supplementary stage is opened to bidders with fresh eligibility rather than restricted to initial regional eligibility (Section 5.1), the separation of routes no longer extends automatically past the clock stage. Two safeguards are then available: excluding a bidder from the supplementary stage in regions where its own package release created or increased excess supply; or, more mildly, charging such a bidder, for any block it acquires there, the regional price implied by its last package price—its package price at release divided pro rata by regional posted prices—rather than the supplementary stage's posted price.

8.7 Other standard options

The remaining options of the standard format carry over unchanged: proxy instructions for reductions at prices above the clock price; a dollar cap on the increment that varies by round; an aggregate reserve price; phases with increasing activity requirements; the grace-round eligibility formula; and the choice between limited and full disclosure.

9. Concluding Remarks

9.1 What the HCA removes from HPB

Relative to HPB, the HCA eliminates bid revival (no bid survives the round in which it is placed, except as processed demand at the current price); current price estimates and their apportionment rule (the package has its own price, and regional prices are set by regional competition); the alternation of provisionally winning status (there are no PWBs); the dependence of the final allocation on the order in which bids were placed (Proposition 2); the SMR pace (prices rise on every product with excess demand in every round); and the costless waiting position that a dormant item bid affords a regional bidder (Section 7.2). It preserves transparent prices; freedom from exposure for package bidders; the possibility that regional bidders collectively outbid a package without coordinating (Theorem 1); and winner determination that requires no optimization. It adds a package premium, which is the price of exposure protection, and the possibility of undersell through package exits, which the supplementary stage addresses. Like HPB, it keeps package bids and regional-level bids separate: neither is ever converted into the other.

In the LLG model the comparison is sharp. The HCA's allocation is a function of the values alone; it is efficient, selling the package exactly when the package bidder's value exceeds the regional bidders' displacement value (Theorem 1); its prices lie in the core, dominate the Vickrey prices, and are Walrasian prices of the hierarchical economy (Theorem 2); with several packages it errs, if at all, only by selling too few, by an amount bounded by the spread of regional values (Theorem 3).

HPB's allocation in the same model depends on the order of bids; the increments it imposes on regional bidders are apportioned shares of package increments, at the scale of the package rather than of the region; its revival of dormant bids can hand a regional bidder a second license it does not want while excluding one that would pay the price (Proposition 2); and a budget-constrained bidder under HPB must solve a portfolio problem that the HCA never poses, because dormant bids encumber its budget while purchasing nothing (Proposition 4). What apportionment and revival were designed to achieve under HPB—keeping regional bids relevant against a package, and permitting regional bidders to displace a package collectively—the HCA achieves structurally, through the reserved units and the sum rule, without either.

9.2 What remains open

The results above are confined to the LLG model and to straightforward bidding. Three questions, each of which the specification is designed to support, remain.

1. **Incentives and the threshold problem.** Proposition 3 shows that straightforward bidding is an ex post equilibrium of the clock stage, but not of the clock stage followed by a supplementary stage at current prices, and identifies the trade-off between undersell and shading in that stage's design. Whether the HCA's incentives are better than HPB's—where an item bidder's incentive is to free-ride on the other item bidders' response to a package, and the remedy, apportionment, is what Proposition 2(c) analyzes—is an equilibrium question, well suited to laboratory experiments of the kind conducted by Brunner et al. (2010) and Kagel et al. (2010). The demand-reduction incentives of uniform-price clock auctions, and the commitment to a route that the no-conversion rule creates, belong to the same question.

2. **Richer environments.** The extensions listed at the end of Section 6.8—multi-unit regional demand, region-specific reserves, bidders of mixed type—appear routine. An intermediate tier (Section 8.2), and the stale-price problem across levels, do not. An exact version of Theorem 3 for an interior m^* requires either the memory variant of Section 6.5 or a supplementary stage that reopens the package; whether either is worth its cost is a design question rather than a theorem.

3. **Pace and undersell in practice.** The number of rounds under the HCA is governed by the increment and the ratio of final to opening prices, as in any clock auction; under HPB, by the number of items and the bid increments. Undersell under the HCA arises only through package exits, and the supplementary stage cures it at the cost identified in Proposition 3. Simulation on realistic value distributions, with realistic increments, would quantify both, and would put numbers on the scale argument of Section 6.6.

References

Ausubel, L. M. (2004). "An Efficient Ascending-Bid Auction for Multiple Objects." *American Economic Review* 94, 1452–1475.

Ausubel, L. M. (2006). "An Efficient Dynamic Auction for Heterogeneous Commodities." *American Economic Review* 96, 602–629.

Ausubel, L. M., and O. V. Baranov (2020). "Core-Selecting Auctions with Incomplete Information." *International Journal of Game Theory* 49, 251–273.

Ausubel, L. M., and P. Cramton (2010). "Virtual Power Plant Auctions." *Utilities Policy* 18, 201–208.

Ausubel, L. M., P. Cramton and P. Milgrom (2006). "The Clock-Proxy Auction: A Practical Combinatorial Auction Design." In P. Cramton, Y. Shoham and R. Steinberg (eds.), *Combinatorial Auctions*, MIT Press.

Ausubel, L. M., P. Cramton, M. Pycia, M. Rostek and M. Weretka (2014). "Demand Reduction and Inefficiency in Multi-Unit Auctions." *Review of Economic Studies* 81, 1366–1400.

Brunner, C., J. K. Goeree, C. A. Holt and J. O. Ledyard (2010). "An Experimental Test of Flexible Combinatorial Spectrum Auction Formats." *American Economic Journal: Microeconomics* 2, 39–57.

Bulow, J., J. Levin and P. Milgrom (2009). "Winning Play in Spectrum Auctions." NBER Working Paper No. 14765. Cambridge, MA: National Bureau of Economic Research.

Bureau of Ocean Energy Management (2024). *Atlantic Wind Lease Sale 10 for Commercial Leasing for Wind Power Development on the U.S. Central Atlantic Outer Continental Shelf—Final Sale Notice.* Federal Register 89(126), 54506–54519, July 1, 2024.

Che, Y.-K., and I. Gale (1998). "Standard Auctions with Financially Constrained Bidders." *Review of Economic Studies* 65, 1–21.

Day, R. W., and P. Milgrom (2008). "Core-Selecting Package Auctions." *International Journal of Game Theory* 36, 393–407.

Federal Communications Commission (2007). *Auction of 700 MHz Band Licenses Scheduled for January 24, 2008: Notice and Filing Requirements, Minimum Opening Bids, Upfront Payments and Other Procedures for Auction 73*, Public Notice DA 07-4171. Washington, DC.

Federal Communications Commission (2019). *Auction 103 Clock Phase Technical Guide*. Washington, DC.

Goeree, J. K., and C. A. Holt (2010). "Hierarchical Package Bidding: A Paper & Pencil Combinatorial Auction." *Games and Economic Behavior* 70, 146–169.

Innovation, Science and Economic Development Canada (2020). *Policy and Licensing Framework for Spectrum in the 3500 MHz Band*, SLPB-001-20, including Annex C, "The Clock Auction Format."

Kagel, J. H., Y. Lien and P. Milgrom (2010). "Ascending Prices and Package Bidding: A Theoretical and Experimental Analysis." *American Economic Journal: Microeconomics* 2, 160–185.

Krishna, V., and R. W. Rosenthal (1996). "Simultaneous Auctions with Synergies." *Games and Economic Behavior* 17, 1–31.

Milgrom, P. (2000). "Putting Auction Theory to Work: The Simultaneous Ascending Auction." *Journal of Political Economy* 108, 245–272.

Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.

Rothkopf, M. H., A. Pekeč and R. M. Harstad (1998). "Computationally Manageable Combinational Auctions." *Management Science* 44, 1131–1147.

Appendix A: Proofs

Proposition 1 is proved first, followed by Theorems 1–3 and then Propositions 2–4.

A.1 Proof of Proposition 1

(a) A regional posted price is the clock price, the price of an applied regional reduction (which lies between the start-of-round and clock prices), or the start-of-round price, all at least p_j^{t-1}. The package posted price is either its clock price, which by Section 4.4 is at least p_0^{t-1}, or $\max\{p_0^{t-1}, \sum_j p_j^t\} \geq p_0^{t-1}$.

(b) If the package ends the round without excess demand, $p_0^t = \max\{p_0^{t-1}, \sum_j p_j^t\} \geq \sum_j p_j^t$, and $\psi^t = \max\{p_0^{t-1} - \sum_j p_j^t, 0\} = \max\{\psi^{t-1} - \sum_j \left(p_j^t - p_j^{t-1}\right), 0\}$. If it ends with excess demand, $p_0^t = P_0^t \geq \sum_j P_j^t \geq \sum_j p_j^t$ in both cases of Section 4.4. Since $p_0^1 = \sum_j m_j$, $\psi_1 = 0$, and the premium can increase only in a round of the second kind. The last statement follows.

(c) For a regional reduction at price point θ that is not applied: the bid was in the queue from θ onward and was re-tested after every application; had $\overline{AD}_j$ exceeded S at any later point it would have been applied. So the product ends without excess demand, and any regional reduction applied after θ would have found $\overline{AD}_j > S$, a contradiction; hence every applied regional reduction has price point at most θ, and the posted price is at most the regional price at θ. For a package reduction that is not applied: it entered the queue in step 2 of processing, before any other bid to change demand, and was re-tested after every application. Hence at no point after step 2 did any region have regional excess demand or did package demand exceed the package supply, since such a state would have been created by an application and the re-test following it would have released the unit. In particular no regional reduction was applied in step 3, so $p_j^t = p_j^{t-1}$ for every j; the package ends without excess demand; and $p_0^t = \max\{p_0^{t-1}, \sum_j p_j^{t-1}\} = p_0^{t-1}$ by (b).

(d) The first statement is the clearing condition $D_0 \leq S_0$. For the second, a regional reduction is applied only while $\overline{AD}_j = D_j + \min\{D_0, S_0\} > S$, which requires $D_j > K$. The third follows from the first two, since winners are the holders of processed demand at the close.

(e) Regional reductions are never applied into excess supply; increases never reduce demand; and the only way $\overline{AD}_j$ can fall below S after having reached it is a package release applied while $D_0 \le S_0$.

(f) The pricing statements are the definitions. For termination: posted prices are bounded below by opening prices and rise by the increment on every product with excess demand; eligibility is non-increasing and is reduced whenever activity falls short of the requirement; with the activity requirement at 100 percent in the final phase of the clock stage, a bidder cannot hold demand whose commitment exceeds any fixed budget indefinitely, and the standard termination argument for clock auctions applies. ■

A.2 Proof of Theorem 1

Proof. We proceed in steps.

Step 1: monotonicity and the order of exits. By construction $p_j(\cdot)$ and $p_0(\cdot)$ are non-decreasing, and $p_0(t) \ge \sum_j p_j(t)$ for all t: this holds at $t = 0$, is preserved by the sum rule, and during an own-increment interval ($D_0 \ge 2$, $H = \varnothing$) $\sum_j p_j$ is constant while p_0 rises. Since a local bidder exits when p_j reaches its value and p_j is non-decreasing, local bidders in region j exit in increasing order of value, and at any time the active local bidders in j are those with values above $p_j(t)$, together with a bidder whose value equals $p_j(t)$ at the instant of its own exit. Likewise global bidders exit in increasing order of value and the active global bidders are those with values above $p_0(t)$. By genericity, no global exit coincides with a local exit.

Step 2: no exit is held. Consider a local bidder's exit at time t, so that $p_j(t) = v_\ell > 0 = p_j(0)$. Since events are finitely many, choose ε so that none occurs on $[t-\varepsilon, t)$; the state is then constant on that interval, and since $p_j(t-\varepsilon) < p_j(t)$ (the exit occurs at the first time p_j reaches v_ℓ), p_j rises on it, so $j \in H$ on it and $j \in H(t^-)$: the exit is applied. Consider a global bidder's exit at t, so that $p_0(t) = V_g > 0$. Choosing ε as before, p_0 rises on $(t-\varepsilon, t)$ with the state constant there, which under the dynamics requires either $D_0 \ge 2$ (own increment) or $H \ne \varnothing$ (the sum rule, since $\sum_j p_j$ rises only when some region is hot); so the exit condition holds at t^- and the exit is applied. Thus the process is a finite sequence of applied exits. Between consecutive events, if $H \ne \varnothing$ some regional price rises at rate one toward the value of the next local bidder to exit, and if $H = \varnothing$ and $D_0 \ge 2$ the package price rises at rate one toward the value of the next global bidder to exit; in either case the next event occurs within finite time, and when $H = \varnothing$ and $D_0 \le 1$ the stage ends. This proves (i).

Step 3: the contest among global bidders. Let t_2 be the first time at which $D_0(t) \le 1$ ($t_2 = 0$ if $|G| \le 1$). For $t < t_2$ the active global bidders number at least two and, by Step 1, include g_1 and the bidder with value $V^{(2)}$; hence $p_0(t) < V^{(2)} < V^{(1)}$ for $t < t_2$, and $p_0(t_2) = V^{(2)}$ if $|G| \ge 2$ (and $p_0(t_2) = 0 = V^{(2)}$ otherwise). In particular g_1 is active at t_2 and, by Step 1, $\sum_j p_j(t_2) \le V^{(2)}$.

Step 4: the package price after the contest. For $t \ge t_2$, $D_0(t) \le 1$, so there is no own-increment interval and the package price is governed by the sum rule alone:

$$p_0(t) = \max\left\{V^{(2)}, \sum_j p_j(t)\right\} \qquad (t \geq t_2, \text{while } g_1 \text{ is active}).$$

Since $V^{(2)} < V^{(1)}$, p_0 can reach $V^{(1)}$ only through the sum: g_1 exits at the first time $t_1 \geq t_2$, if any, at which $\sum_j p_j(t_1) = V^{(1)}$.

Step 5: regional dynamics while the package is demanded. While $D_0(t) = 1$ (that is, on $[t_2, t_1)$, or on $[t_2, \infty)$ if g_1 never exits), region j is hot if and only if $D_j(t) > K$. A region with at most K local bidders is never hot and its price stays at zero. A region with more than K local bidders is hot from time zero (recall that G is non-empty); its price rises at rate one while its local bidders exit in increasing order of value, and it ceases to be hot at the instant its $(K+1)$-th highest-value local bidder exits, which occurs at $p_j = v_j^{(K+1)}$. Thereafter its price is constant and K local bidders—the K highest-value ones—remain active. (The same description applies on $[0, t_2)$, since $\min\{D_0, 1\} = 1$ there as well.) Hence, while $D_0 = 1$, $\sum_j p_j(t)$ is continuous and non-decreasing, reaches Δ in finite time, and never exceeds Δ. Consequently t_1 exists if and only if $\Delta > V^{(1)}$, and in that case $\sum_j p_j(t_1) = V^{(1)} < \Delta$, so at t_1 at least one region is still hot.

Step 6: the case $V^{(1)} > \Delta$. By Step 5, g_1 never exits. Every region with more than K local bidders settles with its K highest-value local bidders active at price $v_j^{(K+1)}$; every other region has all its local bidders active at price zero. Let T be the later of t_2 and the last settling time. At T, $H(T) = \varnothing$ and $D_0(T) = 1$, so the stage ends; it cannot have ended earlier, since before T either some region was hot or $D_0 \geq 2$. By Step 4, $p_0(T) = \max\{V^{(2)}, \Delta\}$. In a region with at least K local bidders, K local winners and the package account for all $K+1$ blocks, so there is no undersell; in a region with fewer than K local bidders, the unsold blocks are reserved units that no local bidder demanded. The allocation has value W_1. This proves (ii).

Step 7: the case $V^{(1)} < \Delta$. By Step 5, g_1 exits at t_1, and by Step 2 the exit is applied, so $D_0 = 0$ from t_1 on and region j is hot if and only if $D_j(t) > K+1$. Regions at t_1 fall into three classes. A region that had already settled has $D_j = K$ active bidders, the K highest; it is not hot, its price stays at $v_j^{(K+1)}$, and $\overline{AD}_j = K < K+1$, so exactly one block is stranded. A region that was still hot at t_1 with exactly $K+1$ active bidders ceases to be hot, its price stays at $p_j(t_1)$, and its $K+1$ blocks are held by its $K+1$ highest-value local bidders. A region that was still hot at t_1 with at least $K+2$ active bidders remains hot; its price rises until its $(K+2)$-th highest-value local bidder exits at $v_j^{(K+2)}$, after which its $K+1$ highest-value local bidders hold its $K+1$ blocks. The stage ends when the last such region settles. In the supplementary stage, each region of the first class offers one block at opening price $v_j^{(K+1)}$; by Step 1 the only non-winning local bidder in that region with value at least $v_j^{(K+1)}$ is the bidder with exactly that value, who demands the block, so demand equals supply at the opening price and that bidder wins it at $v_j^{(K+1)}$. In every region the final winners are therefore the $K+1$ highest-value local bidders, and the allocation has value W_0. This proves (iii).

Step 8: efficiency. By Steps 6 and 7 the format selects the allocation of value W_1 when $V^{(1)} > \Delta$ and of value W_0 when $V^{(1)} < \Delta$, which is the efficient allocation by the computation preceding the theorem. This proves (iv).

Step 9: Vickrey payments when the package is sold. Suppose $V^{(1)} > \Delta$. The Vickrey–Clarke–Groves payment of g_1 is the loss its presence imposes on the others: the efficient value without g_1 is $\max\{V^{(2)} + \sum_j \sum_{i \le K} v_j^{(i)},\ W_0\}$, and the others' value in the efficient allocation with g_1 is $\sum_j \sum_{i \le K} v_j^{(i)}$; the difference is $\max\{V^{(2)}, \Delta\} = p_0^*$. The Vickrey payment of a winning local bidder in region j, with value among the K highest there, is computed as follows. Removing it lowers the displacement value from Δ to $\Delta - v_j^{(K+1)} + v_j^{(K+2)} \le \Delta < V^{(1)}$, so the package is still sold in the efficient allocation without it, and its block passes to the local bidder with value $v_j^{(K+1)}$; the loss imposed on the others is therefore $v_j^{(K+1)} = p_j^*$. This proves (v). ■

A.3 Proof of Theorem 2

Proof. Write T_j for the set of the K highest-value local bidders in region j and T'_j for the set of the $K+1$ highest. By Theorem 1 the outcome is efficient, so for (a) it suffices to show $W(C) \le W(N) - \sum_{i \notin C} u_i$ for every C, since $R + \sum_{i \in C} u_i = W(N) - \sum_{i \notin C} u_i$. For a region j and a set C, let $\text{top}_n(C_j)$ denote the sum of the n largest values among the local bidders of region j that belong to C (zero if there are none).

Case (ii): $V^{(1)} > \Delta$. The winners are g_1 and the members of the T_j; $u_{g_1} = V^{(1)} - p_0^*$ and $u_i = v_i - v_j^{(K+1)}$ for $i \in T_j$. Let $n_j = |T_j \setminus C|$. Then

$$W(N) - \sum_{i \notin C} u_i \ = \ \left[V^{(1)} \text{ if } g_1 \in C,\ p_0^* \text{ otherwise}\right] \ + \ \sum_j \left[\sum_{i \in T_j \cap C} v_i + n_j v_j^{(K+1)} \right].$$

The coalition C can either sell the package to its best global member and give K blocks per region to its best local members, or sell no package and give $K+1$ blocks per region to its best local members. In the first case its value is at most $\max_{g \in C} V_g + \sum_j \text{top}_K(C_j)$, and $\text{top}_K(C_j) \le \sum_{i \in T_j \cap C} v_i + n_j v_j^{(K+1)}$, because the K best members of C_j consist of $T_j \cap C$ together with at most n_j further bidders, each of value at most $v_j^{(K+1)}$; and $\max_{g \in C} V_g$ is $V^{(1)}$ if $g_1 \in C$ and at most $V^{(2)} \le p_0^*$ otherwise. In the second case its value is at most $\sum_j \text{top}_{K+1}(C_j) \le \sum_j \left[\sum_{i \in T_j \cap C} v_i + n_j v_j^{(K+1)}\right] + \Delta$, and $\Delta < V^{(1)}$, $\Delta \le p_0^*$. Either way $W(C)$ is at most the display.

Case (iii): $V^{(1)} < \Delta$. The winners are the members of the T'_j, at prices $p_j(T) \in \left[v_j^{(K+2)}, v_j^{(K+1)}\right]$ with $\sum_j p_j(T) \ge V^{(1)}$ (from the proof of Theorem 1). Let $n'_j = |T'_j \setminus C|$. Then $W(N) - \sum_{i \notin C} u_i = \sum_j \left[\sum_{i \in T'_j \cap C} v_i + n'_j p_j(T)\right]$. If C sells no package, its value is at most $\sum_j \text{top}_{K+1}(C_j) \le \sum_j \left[\sum_{i \in T'_j \cap C}\right.$

$v_i + n'_j v_j^{(K+2)}\Big]$, which does not exceed the display. If C sells the package, its value is at most $V^{(1)} + \sum_j \text{top}_K\left(C_j\right)$. If $n'_j \geq 1$, the K best members of C_j include at most $n'_j - 1$ bidders outside T'_j, each of value at most $v_j^{(K+2)} \leq p_j(T)$, so $\text{top}_K\left(C_j\right) \leq \sum_{i \in T'_j \cap C} v_i + \left(n'_j - 1\right) p_j(T)$; if $n'_j = 0$, then $T'_j \subseteq C$ and $\text{top}_K\left(C_j\right) = \sum_{i \in T'_j} v_i - v_j^{(K+1)} \leq \sum_{i \in T'_j \cap C} v_i - p_j(T)$. In either case $\text{top}_K\left(C_j\right) \leq \sum_{i \in T'_j \cap C} v_i + n'_j p_j(T) - p_j(T)$, and summing, $W(C) \leq V^{(1)} - \sum_j p_j\,(T) + [W(N) - \sum_{i \notin C} u_i] \leq W(N) - \sum_{i \notin C} u_i$. This proves (a).

For (b), compute the Vickrey payments. In case (ii), removing g_1 leaves $\max\{V^{(2)} + \sum_j \text{top}_K\left(L_j\right),\ W_0\}$, so g_1's Vickrey payment is $\max\{V^{(2)}, \Delta\} = p_0^*$; removing $i \in T_j$ leaves $\max\{W_1 - v_i + v_j^{(K+1)},\ W_0 - v_i + v_j^{(K+2)}\}$, so its Vickrey payment is $\max\{v_j^{(K+1)},\ v_j^{(K+2)} - \left(V^{(1)} - \Delta\right)\} = v_j^{(K+1)} = p_j^*$. In case (iii), removing $i \in T'_j$ leaves $\max\{W_0 - v_i + v_j^{(K+2)},\ W_1 - \min\{v_i, v_j^{(K+1)}\} + v_j^{(K+1)}\}$ (if $i \in T_j$ the package allocation replaces i by the $(K+1)$-th bidder; if i is the $(K+1)$-th bidder the package allocation does not use it), so its Vickrey payment is $\max\{v_j^{(K+2)},\ v_j^{(K+1)} - \left(\Delta - V^{(1)}\right)\}$. The HCA price satisfies $p_j(T) \geq v_j^{(K+2)}$ by Theorem 1(iii); and at the instant t_1 at which g_1 exits, $\sum_k p_k\,(t_1) = V^{(1)}$ with $p_k(t_1) \leq v_k^{(K+1)}$ for every k, so that $p_j(t_1) \geq V^{(1)} - \sum_{k \neq j} v_k^{(K+1)} = v_j^{(K+1)} - \left(\Delta - V^{(1)}\right)$, and $p_j(T) \geq p_j(t_1)$. Hence $p_j(T)$ is at least the Vickrey payment.

For (c), in case (ii) the final prices are $p_j^* = v_j^{(K+1)}$ and $p_0^* = \max\{V^{(2)}, \Delta\} \geq \sum_j p_j^*$. At these prices the K members of T_j demand a block, the $(K+1)$-th bidder is indifferent and is assigned none, lower bidders demand none; g_1 demands the package, g_2 is indifferent or demands none; and the seller, facing $p_0^* \geq \sum_j p_j^*$, supplies one package and K blocks per region. Every market clears. Any Walrasian equilibrium in which the package is sold has $p_0 \geq V^{(2)}$ (else a second global bidder demands it), $p_j \geq v_j^{(K+1)}$ (else $K+1$ local bidders demand K blocks) and $p_0 \geq \sum_j p_j \geq \Delta$ (the seller's choice), so (p^*, p_0^*) is the smallest. In case (iii), $p_j(T) \in \left[v_j^{(K+2)}, v_j^{(K+1)}\right]$ and the posted package price is $p_0(T) = \max\{\Pi(T), \sum_j p_j\,(T)\} = \sum_j p_j\,(T)$, since $\Pi(T) \leq V^{(2)} < V^{(1)} \leq \sum_j p_j\,(T)$. At these prices exactly the members of T'_j demand a block in region j (the $(K+1)$-th bidder being indifferent when $p_j(T) = v_j^{(K+1)}$, and assigned the block it wins in the supplementary stage), no global bidder demands the package since $p_0(T) \geq V^{(1)}$, and the seller is indifferent between $m = 0$ and $m = 1$ and supplies $K+1$ blocks per region. Every market clears. The prices are not in general the smallest, because $\sum_j p_j\,(T)$ may exceed $V^{(1)}$ in regions that were still hot at t_1. ■

A.4 Proof of Theorem 3

Proof. (i) The argument of Step 2 in the proof of Theorem 1 applies verbatim, since a local exit occurs while its region's price is rising and a global exit while the package price is rising, and the exit conditions are those under which the respective price can rise.

(ii) As in Step 1, $p_0(t) \geq \sum_j p_j(t)$ throughout, and bidders of each kind exit in increasing order of value, so the active global bidders at T are the $\hat{m}$ highest. Let g be the $\hat{m}$-th highest; it is active at T, so $V^{(\hat{m})} \geq p_0(T) \geq \sum_j p_j(T)$. Since no region is hot at T, $D_j(T) \leq S - \hat{m}$ for every j; as the active local bidders in j are those with values above $p_j(T)$, at most $S - \hat{m}$ local values in j exceed $p_j(T)$, so $p_j(T) \geq v_j^{(S-\hat{m}+1)}$. Summing, $V^{(\hat{m})} \geq \Delta_{\hat{m}}$, and since $V^{(m)} - \Delta_m$ is decreasing in m and never zero, $\hat{m} \leq m^*$.

(iii) Suppose $m^* = S_0$ and, for contradiction, that a global bidder g among the S_0 highest exits, at time t with $p_0(t) = V_g$. If the package price was rising on its own just before t, then $D_0 \geq S_0 + 1$ there, so some active global bidder ranked below S_0 had a value no larger than V_g and would have exited first—a contradiction. Otherwise $p_0(t) = \sum_j p_j(t)$. A region is hot only if it has more than $S - \min\{D_0, S_0\} \geq K$ active local bidders, so its price never exceeds $v_j^{(K+1)}$: $\sum_j p_j(t) \leq \Delta_{S_0} < V^{(S_0)} \leq V_g$, again a contradiction. Hence the S_0 highest global bidders are active at T, every other global bidder exits on an own-increment interval, $\Pi(T) = V^{(S_0+1)}$, and each region ends with exactly K active local bidders at $p_j(T) = v_j^{(K+1)}$; so $p_0(T) = \max\{V^{(S_0+1)}, \Delta_{S_0}\}$. The Vickrey computations are those of Theorem 1(v) with $V^{(2)}$ replaced by $V^{(S_0+1)}$ and Δ by Δ_{S_0}. If $m^* = 0$ then $\hat{m} = 0$ by (ii).

(iv) At T the active local bidders in region j are its $D_j(T)$ highest-value bidders, and they win at $p_j(T)$. The region has $u_j = S - \hat{m} - D_j(T) \geq 0$ stranded blocks. In the supplementary clock for j, which opens at zero, every local bidder that is not already a winner demands one block while the price is below its value; there are no packages, so an exit is applied whenever demand exceeds u_j, which is the case whenever the price is rising, and the clock stops when demand first falls to u_j. By the argument of Step 1, the winners are the u_j highest-value bidders among those not already winning, so the region's $S - \hat{m}$ blocks go to its $S - \hat{m}$ highest-value local bidders. Writing $W(m)$ for the value of the efficient allocation with exactly m packages, $W(m) - W(m-1) = V^{(m)} - \Delta_m$, which gives the equality. For the bound: a global bidder ranked $m \leq m^*$ that is not a winner did not exit on an own-increment interval (on such an interval the exiting bidder is ranked below S_0), so it exited when $\sum_j p_j = V^{(m)}$, and regional prices never exceed $v_j^{(K+1)}$; hence $V^{(m)} - \Delta_m \leq \sum_j \left(v_j^{(K+1)} - v_j^{(S-m+1)}\right) \leq \sum_j \left(v_j^{(K+1)} - v_j^{(S-m^*+1)}\right)$ for every $m \leq m^*$. ■

A.5 Proof of Proposition 2

Proof of (a). Under straightforward bidding every regional reduction is submitted at a price point equal to the bidder's value and, by the argument of Theorem 1 Step 2 transposed to price points, is processed while its region's clock range is being used, so it is applied; and a package reduction submitted at a start-of-round price is the first package reduction of the round and is applied (Section 4.7). The processed demands at the end of each round are therefore determined by the values and the prices alone, and the same path of prices and demands obtains for every processing

order and every tie-break. Local bidders exit at their values exactly; a global bidder maintains demand at the start-of-round package price while the clock-range price does not exceed its value, so its exit price is within one package increment of its value, which gives the stated margin.

Proof of (b). The claim follows from the worked example in Section 6.6 of the main text; every inequality used there is strict, so the set of value profiles is open. ▫ $

Proof of (c). The first statement is the minimum-acceptable-bid rule while the package is provisionally winning; the shortfall after a package raise is δ_0, so a local bidder with value v and high bid $h(a_j)$ cannot respond when $v < h(a_j) + s_j\delta_0 + \delta$, which bounds the gap between its last bid and its value. If region j is frozen at $h(a_j)$, the licenses can overtake a package bid B only if the other regions' bidders can raise their open licenses by $B - \sum_k h(a_k)$ between them, each within its own value; there are profiles with $\sum_k v_k^{(K+1)} > V^{(1)}$ but $\sum_{k\neq j} v_k^{(K+1)} + h(a_j) < B \leq V^{(1)}$, and they form an open set of the stated width. ■

A.6 Proof of Proposition 3

Proof. (a) Fix the others' straightforward strategies. A global bidder g can only choose when to exit. If it exits at $p_0 < V_g$ it earns zero; if it stays while $p_0 \leq V_g$ it either wins at a price $p_0(T) \leq V_g$, earning a non-negative payoff, or exits at V_g and earns zero; staying beyond V_g can only win at a price above V_g. Since the price it pays, $\max\{V^{(2)}, \Delta\}$ or the price at which it is outbid, is determined by the others, straightforward bidding is optimal. A local bidder ℓ in region j likewise chooses when to exit. Under straightforward bidding ℓ is active at T only if $p_j(T) \leq v_\ell$, so its payoff is non-negative. Exiting at $p_j < v_\ell$ earns zero, since without a supplementary stage a bidder that has exited wins nothing. Staying beyond v_ℓ means remaining active while p_j passes v_ℓ; since prices never fall, ℓ then either loses or wins at a price above v_ℓ. So straightforward bidding is optimal for ℓ as well.

(b) Let region j settle at $v_j^{(K+1)}$ at time $t_j < t_1$ under straightforward bidding, so that in case (iii) its $(K+1)$-th bidder, call it ℓ, wins the stranded block in the supplementary stage at $v_j^{(K+1)}$, for a payoff of zero. Let ℓ instead exit at $x \in \left(v_j^{(K+2)}, v_j^{(K+1)}\right)$. The region then settles at x, the regional sum at every later instant is lower by $v_j^{(K+1)} - x$, and if $\Delta - \left(v_j^{(K+1)} - x\right) > V^{(1)}$ the global bidder still exits, region j still strands one block, and ℓ buys it in the supplementary stage at x, for a payoff of $v_j^{(K+1)} - x > 0$. The set of profiles with $V^{(1)} < \Delta - \left(v_j^{(K+1)} - x\right)$ and $t_j < t_1$ is open. ■

A.7 Proof of Proposition 4

Proof. (a) Budget-feasibility is immediate: β holds demand only at price profiles with $\sum_{k\in R_\beta} p_k \leq w$, its reductions are submitted in a region whose price is rising and are therefore applied (the argument of Theorem 1, Step 2), and by Proposition 1(c) a held reduction could in any case cost no

more than the price at which it was submitted. For the allocation: while the global bidder is active, region j has $\overline{AD}_j \geq 3 > S$ whenever $p_j < \theta_j$, so every regional price rises until ℓ_j's reduction at θ_j is processed—applied in full if β has shed or is absent, held at one unit otherwise—and region j's contribution to the package price never exceeds, and eventually reaches, θ_j. The regional sum therefore rises toward Δ, and the global bidder exits (its exit applied, some region being hot) exactly when the sum reaches $V < \Delta$, or wins at $\max\{\Pi, \Delta\}$ when $V > \Delta$. Given the package decision, each region's blocks go to the demands standing at its final price—ℓ_j's inframarginal units, and β's unit where retained—with any block stranded by the global exit in an already-settled region recovered by ℓ_j in the supplementary stage at its exit price. Each such allocation is the efficient one given the package decision and β's retained set, and the argument nowhere uses the realization of (θ, V) or the processing order. When $w \geq \sum_k \theta_k$, the shed rule never triggers, since $\sum_{k \in R_\beta} p_k \leq \sum_{k \in R_\beta} \theta_k$ on every path.

(b)(i) Consider the four-round cycles of the proof of Proposition 2(b), running simultaneously in the two regions of R_β with ℓ_j and β in place of the two unit-demand locals, and tie-breaks under which β takes the open license first and ℓ_j the reserved license. Whenever the package becomes provisionally winning, β's bid on the open license a_j goes dormant and β, following straightforward bidding, takes the reserved license b_j at one increment above its holder; β then holds two live high bids in region j—one dormant, one provisionally winning—and the same in the other region of R_β, four in all, each at the running price level h. A package displacement revives both dormant bids at once, so budget-feasibility requires $4h \lesssim w$ at every such moment, while straightforward bidding drives h toward $\min\{v_\beta, \theta_j\}$; the two are incompatible when w is below roughly $2\sum_k \min\{v_\beta, \theta_k\}$.

(ii) Take $J = 3$, $R_\beta = \{1,2\}$, $v_\beta = 10$, $w = 16$, $\theta_1 = \theta_2 = \theta_3 = 6$, $V = 17 < \Delta = 18$, with the increments and opening bids of Proposition 2(b). The efficient allocation sells no package and gives β one block in each of regions 1 and 2, ℓ_1 and ℓ_2 one block each, and ℓ_3 both blocks of region 3, for total value 44; β's efficient bill is $12 < 16$. Along the cycles, β's four positions reach level $4 = w/4$, at which point any further bid—including a bid superseding one of its own—would carry worst-case liability above 16, and β, being budget-feasible, stops. The local bidders do not stop: ℓ_j values every license at $6 > 4$, is immune to revival (a revived bid hands it a block it wants), and outbids β's provisionally winning positions and, in the course of the continuing cycles, its dormant ones. Once β's provisionally winning bids are outbid, β has no provisionally winning bids and places no new bids, so under the SMR activity rule its eligibility falls to zero once its waivers are exhausted—before ℓ_1 and ℓ_2 have finished superseding its dormant positions—and β cannot re-enter when its encumbered funds are finally released. The item sums rise to 6 per region, the global bidder's last package bid is displaced below $\Delta = 18 > 17$, and the auction ends with ℓ_1, ℓ_2, ℓ_3 winning both blocks of their regions and β winning nothing: total value 36, a loss of $8 = \sum_{k \in R_\beta} (v_\beta - \theta_k)$. Every inequality used is strict, so the construction survives perturbation; and the loss does not shrink with the increments, because what binds is the feasibility constraint—the face value of the live positions—rather than the minimum-bid arithmetic. ■