ASHOKA UNIVERSITY

CENTRE FOR A
People-centric Energy Transition

A Welfare Maximising Auction Structure for the Indian Gas Exchange

October 2026



Authors

Mr. Jolinson Richi, Junior Research Associate, ACPET

Ms. Jahnavi Agrawal, Intern, ACPET, Ashoka University

Dr. Anandajit Goswami, Research Lead, Senior Research Fellow, ACPET

Formatting and Design

Ms. Bipashna Sharma, Communications Associate, ACPET

About Ashoka Centre for a People-centric Energy Transition (ACPET)

The Ashoka Centre for a People-centric Energy Transition (ACPET) is a research-focused, transdisciplinary centre within Ashoka University, India, established to drive a sustainable, equitable, and "people-centric" shift towards net-zero emissions. It bridges the knowledge gap in energy transition by collaborating with industry and government to create scalable solutions, covering areas like renewable energy, policy, and technology. For further information about ACPET, please visit: acpet.ashoka.edu.in

About Ashoka University

Ashoka University is a private university located in Sonipat, Haryana, providing liberal education in the humanities, social sciences, and natural sciences. The University provides a holistic education that is liberal, multidisciplinary, and interdisciplinary, offering a diverse and inclusive space for its students to think deeply and critically, learn across disciplinary boundaries, express themselves creatively, and communicate with meaning to cause impact and change. For further information about ACPET, please visit: ashoka.edu.in

Suggested Citation

Jolinson, R., Agrawal, J., & Goswami, A. (2026). A Welfare Maximising Auction Structure for the Indian Gas Exchange. Ashoka Centre for a People-centric Energy Transition (ACPET), Delhi. Ashoka University, Sonepat, Haryana.
https://doi.org/10.5281/zenodo.23761551

Contents

Abstract 7
1. Literature Review and the Gap 9
1.1 Clearing a market as an optimisation problem 9
1.2 Non-convexities and the existence of prices 9
1.3 The European route: near-equilibrium without side payments 10
1.4 Solution methods 10
1.5 Gas exchanges and the Indian setting 11
1.6 The gap this paper fills 11
2. Introduction 13
2.1 What an exchange is for 13
2.2 The objective that gets you launched 13
2.3 Volume maximisation is one problem, not the only one 14
2.4 The second question: what the published price leaves out 15
2.5 Scope and what follows 15
3. What the Trade Data Shows 17
3.1 Three findings 17
Finding 1: most trades arrive already matched 17
Finding 2: participants are assembling multi-day strips by hand 17
Finding 3: prices differ sharply between delivery points 17
3.2 The diagnosis 18
3.3 What we propose 18
3.4 Why this is safe for the participant 18
3.5 What we are not claiming 19
3.6 Two structural consequences of the narrow question 19
3.7 Contribution and scope 20
4. The Clearing Problem as a Linear Programme 20
4.1 A picture of the primal and the dual 20
4.2 Posing the problem so that the price is not chosen 21
4.3 Where the price comes from: the Lagrangian and the birth of π 22
4.4 Convex and non-convex markets, and why a supporting price exists here 24
4.5 But in reality: when the market is not convex 25
5. The Main Model: MILP and Branch and Cut 26
5.1 Why a MILP, and in what sense it is non-linear 26
5.2 What a block order is 26
5.3 The full model 27
5.4 What the primal solves and what the dual solves 27

5.5 The test each accepted block must pass 28
5.6 Branch and cut: the allocation-and-pricing loop 29
5.7 A worked example 29
5.8 Uniform buyer price, multi-point sellers and opportunity cost 30
5.9 Paradoxical rejection is permitted; branches are rejected, not participants 31
5.10 Numerical tolerances and the solve-time budget 31
5.11 What the engine publishes 32
5.12 The full primal-dual formulation: notation 32
5.13 The primal programme 33
5.14 The dual programme for a fixed block selection 34
5.15 Complementarity: the equilibrium conditions 35
5.16 One MILP for the whole market: the IGX feasible set 36
5.17 Other objectives on the same feasible set 37
6. Key Concepts 39
6.1 The double-sided auction 39
6.2 Uniform pricing 39
6.3 Divisibility and the acceptance fraction 39
6.4 Welfare, and why the price cancels 40
6.5 Primal and dual 40
6.6 KKT conditions 40
6.7 Shadow prices 41
6.8 Convexity, and what it guarantees 41
7. Design Decisions 42
7.1 Bids are ex-hub 43
7.2 Capacity is a hard constraint 43
7.3 Divisible orders by default 43
7.4 One price, no side payments 44
7.5 Price convention: the midpoint of the admissible range 44
7.6 Welfare first, volume second 44
7.7 Reproducibility is a market-rule obligation 45
7.8 Settlement is downstream 45
8. Preliminary Work: Mapping the Cost Stack 46
9. The LP in Practice: A Single Delivery Point 47
9.1 Idea 47
9.2 Method 47
9.3 Results on a hand-checkable market 48
9.4 Results on the synthetic Dahej book 49
9.5 The problem this exposed 50
9.6 Approaches considered and rejected 50

9.7 Ancillary results 51
10. Multiple Delivery Points 52
10.1 Approach 52
10.2 What we deliberately do not model 52
10.3 Method: multi-point orders 52
10.4 Challenges 53
10.5 Scaling 53
11. What Kind of Market Is IGX Today 54
11.1 The intuition of a market 54
11.2 How IGX compares 54
11.3 Defining the current market type 54
11.4 Where it needs to be 54
12. Limitations and Further Work 55
12.1 Delivery-point capacity 55
12.2 Transport pricing 55
12.3 Linepack and multi-period clearing 55
12.4 Sequential sessions and products that are not day-specific 55
12.5 Lots and rounding 56
12.6 Gas quality 56
12.7 Other open items 56
13. Conclusion and Policy Recommendations 57
13.1 Policy recommendations 57
13.2 Toward a market worth entering 58
13.3 What remains 59
14. References 60

List of Figures

Figure 1. The primal and the dual. 20
Figure 2. The clearing flow. 42

List of Tables

Table 1. What the non-convex auction literature assumes, and what the Indian Gas Exchange has. 11
Table 2. Consecutive Daily contracts traded on a single trade date at one delivery point. 17
Table 3. Cleared prices at different delivery points on the same trade date. 17
Table 4. Acceptance rules for buyers and sellers implied by stationarity. 23
Table 5. What the primal solves and what the dual solves. 27
Table 6. Order book for the worked example: one delivery point, two gas days. 30
Table 7. Dual variables and the opportunity costs they measure. 30
Table 8. Notation for the primal-dual formulation. 32
Table 9. Summary of the design decisions. 43
Table 10. Notation for the single-delivery-point programme. 47
Table 11. The hand-checkable market without and with a capacity limit of 150. 49
Table 12. Inputs of the synthetic Dahej book. 49
Table 13. Clearing results for the synthetic Dahej book. 49
Table 14. How IGX compares with the features of a functioning market. 54
Table 15. Policy recommendations, who acts, and how progress is measured. 58

A Welfare Maximising Auction Structure for the Indian Gas Exchange

From Volume Maximisation to Welfare Maximisation, and From Single-Day Orders to Block Participation

Abstract

The literature on clearing auctions with indivisible orders is large, but it was built for electricity: hourly products, coupled zones and a network dispatched by the same operator. It offers no framework for a young, daily, decentralised gas exchange in which every delivery point prices on its own, bids are ex-hub and transport is booked separately. This paper builds that framework for the Indian Gas Exchange.

The Indian Gas Exchange clears by finding the point at which aggregate demand and aggregate supply cross, and taking the quantity traded there as the outcome. That objective was sufficient for a market being started. It is less apt for a market that is deepening, and the exchange's own trade data suggests why.

Across the traded lines reviewed, buy quantity, sell quantity and traded quantity are identical to the MMBtu in approximately 76 per cent of cases, rising to 84 per cent in Day Ahead and Daily contracts and 89 per cent in Intraday. Competitive clearing does not ordinarily produce exact equality line after line. Taken together with runs of consecutive Daily contracts at one delivery point, executed on a single trade date and priced identically to the rupee across every day of the run, the pattern is consistent with requirements that extend across multiple delivery days being agreed privately and registered afterwards.

One likely cause is a missing product. A buyer needing a fixed quantity across six consecutive days cannot express that requirement on the exchange, so they may agree it privately. Improving the clearing objective is therefore necessary, but may not be sufficient on its own.

This paper first builds a linear programme and then extends it to the main mixed-integer model. It poses clearing as constrained welfare maximisation with cleared volume retained as a lexicographic secondary criterion, so that among all allocations achieving maximum gains from trade the engine selects the one moving the most gas. The price at each delivery point and day then falls out of the solution as the Lagrange multiplier on the balance constraint, which is simultaneously the clearing price and the shadow value of one more unit of gas at that point on that day. Because the objective is linear and the feasible set polyhedral, strong duality guarantees that such a supporting price exists, and the Karush-Kuhn-Tucker conditions deliver the acceptance thresholds and the value of scarce delivery capacity as properties of the solution rather than as rules asserted in a rulebook.

The block order removes that guarantee. All-or-none acceptance makes the decision variable binary, the feasible set non-convex, and a supporting uniform price no longer certain. The non-convexity is physically real, originating in single cargoes, take-or-pay minimum guaranteed quantities and indivisible capacity bookings, and cannot be modelled away. The main model sets out the mixed-integer linear programme (MILP) and its branch and cut solution, the rule that no accepted block may be out of the money across its profile, the iterative allocation-and-pricing loop that enforces it, and the per-session measurement of the welfare given up in exchange for a single clean price.

Two findings sit outside the mechanism and are reported as such. Under realistic pipeline constraints, concentration of cleared quantities arises from physical transmission limits rather than from the clearing algorithm. And because natural gas sits outside GST while its transportation sits inside it, state VAT on

the commodity cannot be set against GST on the transmission service, so the cascade falls on the buyer and is invisible in the published price.

The paper then defines the exchange's current market type, explains why it is not yet a market in the economic sense, and sets out policy recommendations to make it one.

Keywords: gas exchange, double-sided auction, uniform price auction, welfare maximisation, linear programming duality, shadow price, block orders, non-convexity, mixed-integer linear programming, branch and cut, delivered cost, GST.

1. Literature Review and the Gap

Market design has studied, in depth, how an auction should decide who trades, how much, and at what price. Almost all of that work, however, was built for electricity. This section reviews it in five strands, moving from the textbook case to the non-convex markets that motivate this paper, and closes by setting out the gap that the paper fills. The rest of the paper follows from that gap.

1.1 Clearing a market as an optimisation problem

The idea that a market price is the dual of a quantity problem goes back to Samuelson (1952), who showed that a spatial price equilibrium can be found by solving a linear programme, with prices appearing as its dual variables. The general result is now textbook material (Bertsimas and Tsitsiklis, 1997; Boyd and Vandenberghe, 2004). When the clearing problem is convex, strong duality holds, the primal and dual optima coincide, and the dual variables form a set of prices at which every accepted order is willing to trade and every rejected order is not. Such prices are a competitive equilibrium in the classical sense.

Auction theory adds the case for a single uniform price. Under pay-as-bid, each participant pays its own bid and so has every reason to shade it below true value. Under a uniform price, the price a participant pays is set by other orders, so bidding close to true value is safe, and the order book keeps carrying information about what the good is worth (Milgrom, 2004). Together these two strands give the textbook design: maximise welfare over divisible orders, and read one uniform price per market from the dual of the balance constraint. This is the design Section 4 sets out for a single gas delivery point.

1.2 Non-convexities and the existence of prices

The textbook design rests on convexity, and convexity fails as soon as some offers are indivisible. Scarf (1994) showed that with indivisibilities a supporting price may not exist at all: no single price can explain an allocation in which some participants are all-or-nothing. Electricity forced the problem into practice, because power plants have start-up costs, minimum stable output and ramping limits, and exchanges let participants express them through block, complex and minimum-income orders.

A large literature followed, and almost all of it tries to get back to, or close to, a convex situation in which shadow prices exist. O'Neill et al. (2005) fix the integer decisions at their welfare-maximising values, solve the linear problem that remains, and use the duals of the fixing constraints as additional prices paid to the participants concerned. Hogan and Ring (2003) and Gribik, Hogan and Pope (2007) look instead for uniform prices that minimise the side payments, or uplift, needed to keep every participant whole; the second leads to convex hull pricing. Garcia-Bertrand, Conejo and Gabriel (2006) and Ruiz, Conejo and Gabriel (2012) use primal-dual formulations that minimise the duality gap introduced by integer decisions, adding conditions so that producers recover their costs. Liberopoulos and Andrianesis (2016) compare these schemes and show that each trades off efficiency, price simplicity and the size of side payments in a different way.

These approaches differ in detail, but they share one feature. To keep a single price, they pay some participants outside it. This is broadly the North American route.

1.3 The European route: near-equilibrium without side payments

European power exchanges took a different route. Their day-ahead auction, now run across the continent through the EUPHEMIA algorithm, admits block orders but pays no side payments (NEMO Committee, 2020). The rules are asymmetric. An accepted block must never lose money at the published prices. A block that would have made money may still be rejected, and is then reported as a paradoxically rejected block. Every divisible order, and every network constraint, must be at equilibrium. The result is a near-equilibrium with one clean price per area and hour, where the only deviation from a perfect equilibrium is the opportunity cost borne by paradoxically rejected blocks.

Meeus, Verhaegen and Belmans (2009) measured what these restrictions cost in welfare. Martin, Müller and Pokutta (2014) formulated the search for strict linear prices in such markets. Madani and Van Vyve (2015a) gave an efficient mixed-integer formulation for the European auction, with a decomposition that adds locally strengthened cuts to remove price-inconsistent block selections. O'Neill, Sotkiewicz and Rothkopf (2007) showed that if paradoxically accepted blocks were allowed and compensated, there would always be enough welfare to pay for it.

The work closest to this paper is the primal-dual framework of Madani and Van Vyve (2015b). Their key step is to write the primal constraints, the dual constraints and one inequality between the two objectives into a single mixed-integer programme. In the notation of this paper, for welfare W and a dual objective made of order surpluses s and network or capacity values, they impose:

$$W = \sum_{o} \sigma_o \, p_o \, q_o \, x_o \;\geq\; \sum_{o} s_o + \sum_{m} w_m \, v_m \tag{1}$$

Weak duality makes the reverse inequality automatic, so this one constraint forces the two objectives to be equal. That in turn enforces every complementarity condition of the convex part of the market without writing any of them down, and without introducing auxiliary variables. Big-M terms then allow a block to be paradoxically rejected, with a variable that bounds its opportunity cost, while forbidding it to be paradoxically accepted. The framework has two consequences that matter directly for this paper. First, it gives the first tractable MILP formulations of objectives other than welfare, in particular the maximisation of traded volume and the minimisation of opportunity costs of paradoxically rejected blocks. In a non-convex market, maximising welfare and maximising volume are not the same problem.

1.4 Solution methods

The methods behind these formulations are standard in integer programming. Branch and bound, introduced by Land and Doig (1960), relaxes integer variables, solves the resulting linear programme and branches on any variable that comes out fractional. Branch and cut, developed by Padberg and Rinaldi (1991), adds valid inequalities, or cuts, inside the search tree to remove fractional or otherwise unwanted solutions. Benders decomposition (Benders, 1962) splits a problem into a master problem over the integer decisions and a subproblem over the continuous ones, and combinatorial Benders cuts (Codato and Fischetti, 2006) are its yes-or-no version, cutting off one combination of integer decisions at a time. The EUPHEMIA algorithm and the decomposition of Madani and Van Vyve (2015a) both belong to this family. Section 5 uses the same family for gas.

1.5 Gas exchanges and the Indian setting

Gas has received far less attention. The literature on European traded gas hubs, notably the Oxford Institute for Energy Studies series by Heather, is largely institutional. It describes how hubs mature, how liquidity and churn are measured, and which regulatory steps helped. It rarely writes down a clearing model, because most European spot gas trades continuously on screens rather than in a sealed auction, and because gas pipelines are booked separately from the commodity rather than dispatched with it.

For India, the literature is mostly about pricing policy: administered prices for domestic gas, ceiling prices for difficult fields, and the indexation of imported LNG. The Indian Gas Exchange began trading in 2020 under the PNGRB (Gas Exchange) Regulations, 2020. We found no published work that models how IGX clears, that tests its outcomes against its own trade data, or that asks whether it yet functions as a market.

1.6 The gap this paper fills

The literature therefore supplies the right tools, but it built them for a different market. Table 1 sets out the differences.

Table 1. What the non-convex auction literature assumes, and what the Indian Gas Exchange has.

	What the literature assumes	**What IGX actually has**
Setting	Hourly power markets; network and energy dispatched by one operator, so zones are coupled	Daily gas delivery at 20+ independent points; ex-hub prices; transport booked separately with the transporter
Objective	Welfare maximisation is taken as given; volume appears only as an alternative objective to study	A clearing rule that finds the crossing point and maximises traded volume
Products	Hourly orders and short power blocks	Ten tenors from Intraday to 6 Month, and no way to bid for a multi-day package
Price structure	One zonal price for all participants in a zone	One uniform price for buyers at each point; sellers who can deliver at several points face several prices
Evidence	Mature, liquid exchanges	A thin exchange where trade data suggests many deals are agreed privately and only registered

No existing framework combines, in a decentralised daily gas setting, four things: a uniform price at each delivery point for buyers; sellers who may be eligible at several points; a rule that puts welfare first and volume second; and multi-day block orders priced without side payments. Nor does any study ask why such an exchange is not yet a market, or what would make it one.

This paper fills that gap. It adapts the European near-equilibrium approach, and in particular the primal-dual framework of Madani and Van Vyve (2015b), to the Indian gas setting. It replaces volume maximisation with welfare first and volume second, and shows how the primal and dual together keep volume high while putting a rupee value on every source of lost opportunity. It uses IGX's own trade data to argue that a missing block product, and not only the objective, keeps price discovery off the

exchange. And it defines the exchange's current market type and the policy steps that would move it towards a competitive market.

The paper asks three questions. How should a gas exchange like IGX clear, so that its price is derived rather than chosen? How can it offer multi-day block orders and still publish one clean price per point and day? And what would turn IGX from a settlement platform into a competitive market? Section 2 introduces the problem in the Indian context.

2. Introduction

2.1 What an exchange is for

An exchange exists to answer three questions about a delivery period: who trades, how much, and at what price. It is worth being precise about why those questions need a central answer at all, because gas was traded in India for decades without one.

Bilateral negotiation answers all three, but privately and pairwise. Two parties who find each other agree on a quantity and a price, and the agreement tells nobody else anything. A third buyer who would have paid more never learns that gas was available. A fourth seller who could have supplied more cheaply never learns that demand existed. The gains from trade that go unrealised in that situation are not recorded anywhere, which is precisely what makes them easy to ignore.

An exchange does two things bilateral trade cannot. It collects all the willingness to pay and all the willingness to supply for a delivery period into one place, so that the allocation can be made against the whole of both sides rather than against whoever happened to pick up the telephone. And it produces a public number that everyone can see and that everyone who traded actually faced. The first is an efficiency argument and the second is an information argument, and the second is the one that compounds, because a price series that means something becomes the reference against which contracts are indexed, investments are appraised and regulation is set.

India's target of lifting natural gas from a low single-digit share of primary energy to roughly 15 per cent by 2030 depends on both. It is met in molecules delivered, and the molecule arrives only when an industrial user weighing gas against coal finds that gas wins. That comparison is not made at the exchange price. It is made at the burner tip, against a delivered cost that includes zonal transmission tariff, marketing margin, and a tax stack we return to in Section 8. Transmission tariff and tax are therefore not settlement details. They are the margin on which the target turns, and an exchange is the natural place to make them visible because it observes price at the moment of formation.

The Indian Gas Exchange was established to do this work. Section 1 showed that the market design literature offers the right tools, but not for a market like this one. This section argues that the way IGX currently poses the clearing question limits its ability to do this work, and that the remedy is a change of objective together with a new order type, both available without any change to the physical or regulatory arrangements around the exchange.

2.2 The objective that gets you launched

IGX runs a double-sided auction under uniform pricing. Buy orders are ranked from the highest price downwards and sell orders from the lowest price upwards, the two are accumulated into cumulative demand and supply curves, and the algorithm finds the point at which the curves cross. That intersection defines the Market Clearing Price and the Market Clearing Volume. Every successful buyer pays the same price and every successful seller receives it. Ties at the clearing price are broken by price-time priority. Longer contracts clear in open auction windows where participants can watch the book and adjust; Daily and Intraday contracts match continuously.

The objective this pursues is the maximum aggregate quantity of gas that can be traded.

That is a reasonable objective for a new exchange, and it should be said plainly that it was the right choice at launch. A venue with no history has to demonstrate that it can move gas at all. Volume is the metric a regulator understands, the metric a prospective member asks about, and the metric that determines whether the exchange survives its first years. Maximising it directly is a defensible bootstrapping decision.

It is not, however, the same objective as maximising the value the market creates, and the two can select different allocations. Consider two candidate outcomes for the same order book. In the first, a buyer valuing gas at Rs 2,000 per MMBtu is supplied by a seller whose floor is Rs 600, and 100 MMBtu moves. In the second, that same buyer is supplied, and in addition a buyer valuing gas at Rs 610 is supplied by a seller whose floor is Rs 605, moving a further 100 MMBtu. The second clears twice the volume. The value it adds over the first is Rs 500 in total, on a trade whose surplus per unit is five rupees. Both trades are individually rational and neither should be prohibited, but an objective that treats the second hundred as equal in worth to the first has stopped measuring anything about value. Where capacity is scarce, and Section 9 shows that at Dahej it is scarce by a factor of more than five, the two objectives do not merely differ in emphasis. Volume maximisation fills the pipe with the trades that are easiest to match rather than the trades that are worth the most, and the shadow value of the capacity those trades consume is never computed, because nothing in the calculation asks what a unit of throughput is worth.

The direction of travel this paper proposes is therefore from volume maximisation to welfare maximisation. Welfare here means buyers' declared value less sellers' declared cost, summed over what is actually accepted.

Volume is not discarded in the move. Gas is a quantity-driven market and an exchange that appeared indifferent to how much gas it moved would not deserve a hearing. Volume is demoted from objective to tie-break. The engine first finds the maximum achievable welfare, then, among all allocations achieving it, selects the one that clears the most gas. This is a lexicographic rule, implemented in the staged algorithm of Section 9.2 and recorded as a design decision in Section 7.6. It matters more than it looks. Linear programmes routinely have multiple optimal solutions, particularly when orders share limit prices, and an engine that leaves the choice among them to the solver's internal pivoting order is neither reproducible nor defensible. Welfare first, quantity second, time priority third is a complete and reproducible ordering.

2.3 Volume maximisation is one problem, not the only one

The objective is not the whole story. IGX's own trade data points to a second issue that may matter as much: the products on offer may not match how participants need to buy gas. Section 3 sets out the evidence in full.

For 76 per cent of traded lines, buy quantity, sell quantity and traded quantity are identical to the MMBtu, rising to 84 per cent in Day Ahead and Daily contracts and 89 per cent in Intraday. When many buyers meet many sellers, aggregate demand and aggregate supply rarely land on the same number, and somebody is usually left partially unfilled. Alongside that, single trade dates repeatedly produce runs of consecutive Daily contracts at one delivery point, priced identically to the rupee across

the whole run, with quantity shapes that are not flat. These patterns are consistent with packages agreed between two counterparties before the auction and registered on it afterwards.

Other readings are possible. A thin book with one buyer and one seller would also produce equal quantities, and participants may prefer to register trades on the exchange for credit or reporting reasons. We do not claim to prove intent from this data. We treat it as evidence worth explaining.

One explanation fits well. A buyer who needs 54,000 MMBtu per day for six consecutive days has no way to express that requirement on the exchange, because no strip or block product exists. Submitting six separate Daily orders risks being filled on four of them, which leaves a broken requirement to cover elsewhere. A private package removes that risk. If this is part of the cause, a block bid gives such a buyer a way to bring the package into the auction, where it competes against every other order.

The paper therefore treats two changes as complementary. Welfare maximisation improves how orders are cleared. Block bids may bring more orders to be cleared. Improving the objective is worth doing, not least because it is the foundation on which the missing product can be built.

2.4 The second question: what the published price leaves out

The paper has a second part, which stands on its own and which we include because the first part supplies its benchmark.

Once the engine produces a price that can be trusted as the market's own valuation of gas at a delivery point, it becomes possible to ask how much of what a buyer actually pays was formed by the market and how much was added afterwards. Section 8 traces that stack from wellhead or FOB through terminal, grid and distribution network to the burner tip, and the answer is that the exchange price is a minority of the total.

One component of that stack is structural rather than incidental. Natural gas sits outside GST and is taxed under state VAT at rates between 5 per cent and roughly 14.5 per cent, while its transportation is a service inside GST at 18 per cent. The two regimes do not interact, so tax paid on the gas cannot be set against tax paid on moving it, and the cascade compounds along the route for any buyer drawing gas across state boundaries. This is the cascade GST was designed to eliminate everywhere else in the economy. The burden falls hardest, proportionally, on the buyer taking modest volumes over long distances, which is exactly the buyer the exchange most needs to attract and exactly the one for whom gas loses against coal at the burner.

The inefficiency here is at the tax interface, not in the clearing algorithm. The engine allocates correctly at the delivery point and the tax stack then undoes that downstream. Bringing gas within GST would remove a deterrent no clearing rule can reach. Until then, the exchange can publish the decomposition alongside the clearing price, which requires no policy change and makes the burden visible to those in a position to act on it.

2.5 Scope and what follows

Three things are deliberately outside this paper. The engine prices ex-hub throughout, so tariff and tax never enter the objective and settlement is a separate downstream pass. No pipeline flows between delivery points are modelled, because the exchange neither dispatches gas nor holds transmission

capacity rights. And continuous matching under price-time priority, which IGX runs for intraday contracts, produces a sequence of bilateral trades rather than one clearing price and is not addressed.

Section 3 sets out the trade-data findings and what we propose in response. Section 4 builds the clearing problem as a linear programme, starting from a picture of the primal and the dual, and shows where the price comes from. Section 4 ends with what happens in reality, when orders are indivisible and the market is no longer convex. Section 5 presents the main model for that case, a mixed-integer linear programme solved by branch and cut. Section 6 restates the key concepts, Section 7 records the design decisions as a single clearing flow, and Section 8 maps the delivered cost stack. Sections 9 and 10 run the linear programme on order books at one and at many delivery points. Section 11 defines what kind of market IGX is today, Section 12 sets out the limitations, and Section 13 concludes with policy recommendations.

3. What the Trade Data Shows

The argument of Section 2.3 rests on the exchange's own published trade data. The observations below come from IGX published trade data over a recent one-month period, covering 193 contract lines across 14 delivery points.

3.1 Three findings

Finding 1: most trades arrive already matched

For 76 per cent of all traded lines, buy quantity, sell quantity and traded quantity are identical to the MMBtu. For Day Ahead the figure is 84 per cent, for Daily 84 per cent, and for Intraday 89 per cent.

Competitive auction clearing does not ordinarily produce this. When many buyers meet many sellers, aggregate demand and aggregate supply do not land on the same number, and one side is left partially unfilled. Exact equality, line after line, means one buyer and one seller arrived with a quantity already agreed between them.

Finding 2: participants are assembling multi-day strips by hand

Repeatedly, a single trade date produces a run of consecutive Daily contracts at the same delivery point, the same price, and a deliberate quantity shape:

Table 2. Consecutive Daily contracts traded on a single trade date at one delivery point.

Trade date	Point	Days	Price	Quantities
04 Aug	Dahej	5 consecutive	Rs 2,056	54,400 × 4, then 40,000
05 Aug	Dahej	8 consecutive	Rs 2,036	30,000 flat
06 Aug	Hazira	7 consecutive	Rs 1,980	15,000, 10,000, then 25,000 × 5
07 Aug	Mhaskal (Non-HPHT)	5 consecutive	Rs 2,046	18,000 flat

Identical prices across every day of a run, executed on one trade date, with shapes that are not flat. A price that does not move by a single rupee across eight separate contracts is not eight independent clearings. It is one negotiated package, registered as eight contracts because there is no other way to record it.

Finding 3: prices differ sharply between delivery points

On the same trade date, for the same product, at different delivery points:

Table 3. Cleared prices at different delivery points on the same trade date.

Date	Points	Cleared prices
03 Aug	Bokaro / Dahej	Rs 1,750 / Rs 2,047
06 Aug	Hazira-ONGC / Hazira	Rs 500 / Rs 1,980
07 Aug	Hazira-ONGC / Mhaskal (Non-HPHT)	Rs 1,000 / Rs 2,078
10 Aug	Hazira-ONGC / Chhara	Rs 1,250 / Rs 1,955

Gaps of Rs 300 to Rs 1,480, persisting day after day. If gas moved freely between these points, arbitrage would close the gaps within a session. It does not. Delivery points are separate markets rather than interchangeable locations within a hub, and Section 10.1 draws out the consequences for the multi-point formulation.

3.2 The diagnosis

Put the first two findings together.

A buyer needing 54,000 MMBtu per day for six consecutive days has no way to express that on the exchange. There is no strip product. Their options are to submit six separate Daily orders and risk being filled on four of the six, leaving a broken requirement they must cover elsewhere; or to telephone a seller, agree the package privately, and register six contracts at the agreed price.

They choose the second, and the data shows them doing it. The choice is not between the exchange and a bilateral deal. It is between a bilateral deal and nothing, because the product they want does not exist on the exchange.

The consequence is that the exchange performs settlement and record-keeping for prices formed elsewhere. Gas trades, but price discovery does not happen. The published price is a report of a private negotiation rather than the outcome of competition.

This is a missing product, not a defective algorithm. No improvement to clearing mathematics helps if the orders never arrive.

3.3 What we propose

A block order, or strip. One order covering an arbitrary set of delivery days, with a chosen quantity on each day, accepted or rejected as a single package.

The buyer above submits one order: 54,000 per day, six named days, one limit price, all six days or none. Not six orders. One.

The days need not be consecutive or uniform. Fourteen days, then two, then three, is one order. So is every weekday in March at 20,000 with weekends at 5,000. The engine does not care about the shape.

What this changes is where the package is priced. The package that would have been negotiated privately is instead submitted to the auction, where it competes against every other order at those delivery points and days. If a Daily buyer values gas on day three more highly, they win it and the block is rejected. If the block creates more total value, it clears. The price is discovered rather than agreed.

3.4 Why this is safe for the participant

A participant will only bring a package to the auction if the auction cannot harm them. Four properties, each derived in Section 5, provide that.

- **They are never partially filled.** A block is all or nothing. A participant who would have to return to the bilateral market for the missing days is never put in that position.

- **They are never accepted at a loss.** Acceptance requires the package to be in the money across its whole profile. A block that would leave them worse off than not trading cannot be accepted.
- **Rejection costs them nothing.** A rejected block is reported and returned intact. They are free to negotiate privately exactly as before. Nothing is lost by trying the exchange first.
- **The price can be checked.** The published price follows a written rule, is reproducible, and can be re-derived months later. A participant leaving a negotiated price for an auction price can defend that decision internally.

3.5 What we are not claiming

- **Not that existing clearing is broken.** Fifteen sessions in the data had buyers and sellers both present and no trade, with best bids far below best asks. Those are genuine cases of no overlap. This engine would not have cleared them either. The existing method is exact and deterministic: the same orders always produce the same result, with no approximation and no guesswork. Any critique that assumes the algorithm is sloppy is aimed at something that is not there. The gap is that the question being posed is narrower than the market requires.
- **Not that bilateral trading should stop.** Participants keep every option they have today. The proposal adds one.
- **Not that the engine creates liquidity by itself.** If both parties still meet privately before the auction opens, nothing changes. The engine makes the product possible. Whether participants use it is a matter for them and for IGX.
- **Not a change to what is traded.** Delivery obligations, settlement, credit and physical delivery are all unaffected. The block order is a way of expressing demand that already exists.

3.6 Two structural consequences of the narrow question

Beyond the missing product, posing clearing as an intersection-finding exercise has two further effects.

Orders are not partially fillable. An order is taken or it is not, and a participant cannot be filled to 60 per cent of what they asked for. This forces the cleared quantity onto whichever discrete combination of whole orders happens to balance, which is rarely the quantity that maximises gains from trade. It also means the marginal order, the one that should set the price, has no way to express its marginality, and as Section 4.3 showed, it is precisely the partially filled order that pins the price down.

Capacity is a boundary rather than a variable. The auction picks an allocation that makes economic sense, and the system then checks whether it fits within physical delivery limits. If it does not, matched trades are cancelled. The mechanism can therefore choose a result that looks correct on paper but cannot be delivered, and only afterwards trim it back. A limit applied after the fact can invalidate an allocation but cannot shape it, and where offered volume far exceeds throughput it is capacity rather than the crossing of bids and asks that rations the market. A price set without reference to the binding constraint is the wrong price, and nothing in the current calculation produces a value for the constraint that is doing the rationing.

3.7 Contribution and scope

The paper restates clearing as constrained welfare maximisation with cleared volume as a lexicographic secondary criterion, and recovers price as the shadow value of the balance constraint. This yields divisible acceptance, capacity inside the optimisation, a rupee value for scarce capacity, and acceptance conditions that are derived rather than asserted. It then introduces the block order, sets out the non-convex machinery required to price it, and finally takes the resulting price as a benchmark against which to decompose what tariff and tax add downstream.

The paper first builds the clearing problem as a linear programme (Section 4), then extends it to block orders as a MILP solved by branch and cut (Section 5), and then runs the linear programme at one and at many delivery points (Sections 9 and 10). The engine prices ex-hub throughout, and tariff and tax never enter the objective.

4. The Clearing Problem as a Linear Programme

We now write the clearing problem down. This section treats the case in which every order is divisible, so the market is convex and the problem is a linear programme. It starts with a picture of the two halves of that programme, the primal and the dual, then derives each piece in turn.

4.1 A picture of the primal and the dual

Every linear programme comes in a pair. The primal asks the physical question: who trades, and how much gas moves? The dual asks the economic question: what is each constraint worth, and what prices support that physical answer? Figure 1 shows both for the small market used in Section 9.3, with three buyers bidding Rs 1,000, Rs 800 and Rs 500 and three sellers asking Rs 300, Rs 600 and Rs 900, for 100 MMBtu each, at a point whose capacity is 150 MMBtu.

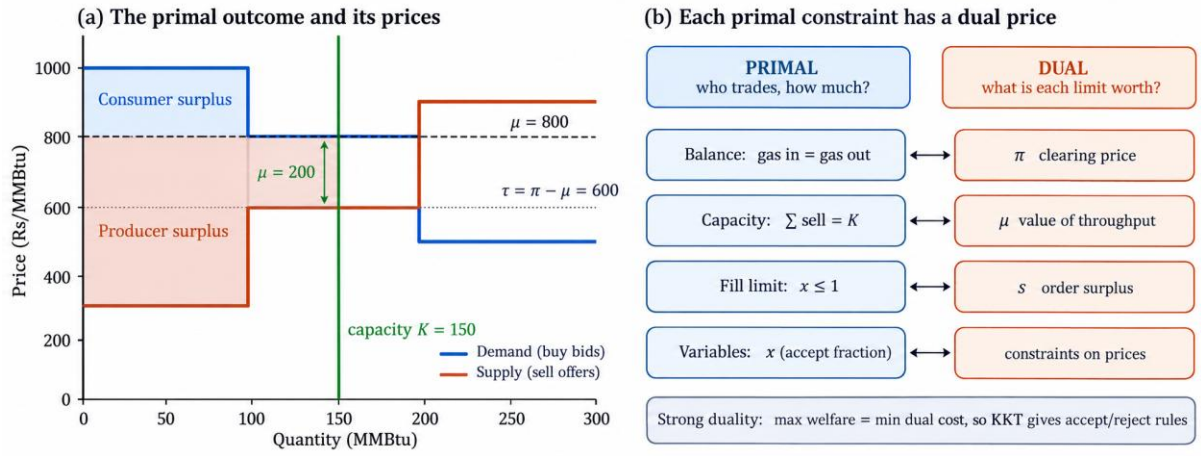


Figure 1. The primal and the dual.
Panel (a) shows the primal outcome on the supply and demand curves; panel (b) shows which dual price belongs to which primal constraint.

Panel (a) is the primal outcome. Capacity binds at 150 MMBtu, so the market clears at the top of the demand curve and the bottom of the supply curve. The clearing price $\pi = 800$ is where the marginal buyer, the one only half filled, sits. The sellers' threshold $\tau = 600$ sits lower, and the gap between them

is μ = 200, the value of one more unit of capacity. The shaded areas are consumer and producer surplus, and their sum is the welfare the primal maximises.

Panel (b) shows where those prices come from. Each constraint of the primal has a partner variable in the dual. The balance constraint's partner is the clearing price π. The capacity constraint's partner is μ. Each order's fill limit has a partner s, that order's surplus. Strong duality says that, at the optimum, the maximum welfare of the primal equals the minimum cost of the dual, and the conditions that link the two, the KKT conditions, are exactly the rules for accepting, rejecting or partly filling an order. The rest of this section derives each part of the picture.

4.2 Posing the problem so that the price is not chosen

We start with the simplest case: one delivery point, one delivery period, and every order divisible. Everything later in the paper, including the MILP of Section 5, is built on this programme.

Let B be the set of buy orders and S the set of sell orders at one delivery point for one delivery period. Buy order b carries a limit price p_b, the most it will pay per MMBtu, and a requested volume q_b. Sell order s carries a floor price p_s and an available volume q_s. Let K be the throughput available at the point for the period.

The only quantities the engine chooses are acceptance fractions. For each order it selects a value between 0 and 1, where 1 is a full fill, 0 a rejection, and 0.22 a fill of 22 per cent. The gas that moves is the requested quantity multiplied by the fraction, not the fraction itself.

The programme is:

$$\max_{x} \; W = \sum_{b \in B} p_b \, q_b \, x_b - \sum_{s \in S} p_s \, q_s \, x_s \tag{2}$$

subject to

$$\sum_{b \in B} q_b \, x_b - \sum_{s \in S} q_s \, x_s = 0 \qquad [\pi] \tag{3}$$

$$\sum_{s \in S} q_s \, x_s \le K \qquad [\mu] \tag{4}$$

$$0 \le x_b \le 1, 0 \le x_s \le 1 \qquad \forall b \in B, \; s \in S \tag{5}$$

In words: choose how much of each order to fill so that buyer value less seller cost is as large as possible, while gas in equals gas out, the point is not overloaded, and nobody receives more than they asked for.

The first thing to notice is what is absent. There is no price variable. The objective contains limit prices and floor prices, which are declarations made by participants, but it contains no price at which anything settles. This is not an oversight, and it is the structural fact on which everything else rests.

To see why, write out the surplus the participants actually receive at some settlement price. A buyer accepted at a given fraction gains the difference between their limit and the price on every unit they

take, and a seller gains the difference the other way. Total surplus is consumer surplus plus producer surplus:

$$CS + PS = \sum_{b} (p_b - \pi)\, q_b\, x_b + \sum_{s} (\pi - p_s)\, q_s\, x_s \tag{6}$$

Expanding and collecting the terms in π:

$$CS + PS = \left[\sum_{b} p_b q_b x_b - \sum_{s} p_s q_s x_s\right] - \pi\left[\sum_{b} q_b x_b - \sum_{s} q_s x_s\right] \tag{7}$$

The first bracket is W. The second bracket is exactly the balance constraint, which is zero for every allocation the engine is allowed to choose. The price cancels. Total surplus equals W regardless of what π is.

This is the π-cancellation, and its consequence surprises people. The money passing from buyer to seller is a transfer, not a gain. The settlement price determines who captures the surplus; it has no effect whatever on how large the surplus is. It follows that the welfare-maximising allocation can be computed without knowing the price, which is what makes it legitimate to write a programme with no price variable in it. Section 9.3 demonstrates this on a hand-checkable market: recomputing the surplus split at three different prices across the admissible range moves surplus between the two sides and leaves the total unchanged at every price.

This is also the central defence of the design. When a participant or a regulator asks how a price arose, the answer is a derivation rather than a discretion. Gas was allocated so as to create the most value, and this is what gas turned out to be worth at that delivery point on that day under that allocation. The exchange never picks a number, and therefore cannot have picked a wrong one.

4.3 Where the price comes from: the Lagrangian and the birth of π

The price is not absent from the problem. It is in the dual.

The Lagrangian is a single expression that combines the objective with every constraint, each constraint multiplied by its own shadow value. Its purpose is mechanical: differentiating it tells us what every shadow value must equal. That is how the prices come out. The signs are chosen so that violating a constraint reduces the total. The multiplier on balance is π and the multiplier on capacity is μ:

$$L(x, \pi, \mu) = \sum_{b} p_b q_b x_b - \sum_{s} p_s q_s x_s$$
$$- \pi\left[\sum_{b} q_b x_b - \sum_{s} q_s x_s\right] - \mu\left[\sum_{s} q_s x_s - K\right] \tag{8}$$

Differentiate with respect to each acceptance fraction:

$$\frac{\partial L}{\partial x_b} = q_b\,(p_b - \pi), \frac{\partial L}{\partial x_s} = q_s\,(\pi - p_s - \mu) \tag{9}$$

Each derivative is that order's surplus per unit, multiplied by its size. **This is where π is created.** It is not introduced as a price and then justified afterwards. It appears as the multiplier the Lagrangian requires on the balance constraint, a purely mathematical object with no economic content attached to it at birth, and the derivative above then shows that the quantity it must be compared against is the participant's own limit price. The object was constructed as a multiplier and turns out to behave as a price.

The Karush-Kuhn-Tucker conditions make this precise. At the optimum of a linear programme, four conditions hold together, and satisfying all four is proof that no better allocation exists anywhere in the feasible set. Primal feasibility says the answer is legal. Dual feasibility says the multipliers carry the right signs, with μ non-negative and π free. Complementary slackness says a constraint that is not binding carries zero shadow value: if a delivery point is below capacity, delivery room there is free. Stationarity says no small nudge to any variable improves the total, and it is obtained by setting the derivatives above to zero.

Applying stationarity to a variable bounded between 0 and 1 gives the acceptance rules directly. If the buyer's limit exceeds π the derivative is positive, so welfare rises by accepting more, and the fraction is pushed to its upper bound: the order is fully accepted. If the limit is below π the fraction is pushed to zero: rejected. Only if the limit equals π can the fraction rest strictly between the bounds, and that order is the marginal one, the one that sets the price. Sellers behave identically, measured against an effective threshold $\tau = \pi - \mu$ rather than against their floor alone.

Table 4. Acceptance rules for buyers and sellers implied by stationarity.

Buyers	**Sellers**	**Outcome**
limit above π	floor below $\tau = \pi - \mu$	fully accepted
limit below π	floor above τ	rejected
limit equal to π	floor equal to τ	partially filled, marginal

These rules are not asserted by the engine. They are properties any optimal solution must have. A procedural engine has to code them by hand and can code them wrongly; the optimisation gets them right by construction. The asymmetry between the two columns is structural rather than a bias. Buyers appear only in the balance constraint, while sellers appear in balance and in capacity, so every unit a seller supplies consumes scarce delivery room and their effective threshold falls by the value of what it consumes. When capacity is slack, μ is zero and the asymmetry vanishes. This result follows from the capacity constraint being written on the delivery-in side; if a limit were later applied to withdrawals, τ would take a different form and must be re-derived from stationarity rather than carried across.

The same multiplier has a second reading, and this is why it carries two names. Perturb the right-hand side of the balance constraint, as though one more MMBtu appeared at the point for free, and the optimal welfare changes at rate:

$$\frac{\partial W^*}{\partial \varepsilon} = \pi \tag{10}$$

That free unit would go to the highest-valued buyer not currently served, and be worth exactly that buyer's bid, and that buyer sits precisely at the margin of the market. The value of one more unit and the price at which the market clears are the same number, reached from two directions. This is why the

multiplier is called a shadow price and the clearing price interchangeably. They are not two quantities that happen to coincide. They are one quantity with two interpretations, and the identity between them is what makes the published price economically meaningful rather than merely conventional.

The capacity multiplier μ reads the same way, and measures the rupee value per day of one more MMBtu of throughput at the point. Where the constraint is slack it is zero and capacity is free. Where it binds, μ multiplied by the volume shortfall reproduces the welfare loss exactly, and that identity closing numerically is what it means for a number to be a shadow price. Section 9.3 checks it.

The primal and the dual are therefore answering two halves of one question. The primal asks the physical question, who trades and how much gas moves, and returns the acceptance fractions. The dual asks the economic question, what is each constraint worth and what prices support that physical answer, and returns π and μ. One solve produces both, which is the practical point: the exchange is not running an allocation calculation and then a separate pricing calculation that might disagree with it. A common misreading is that the dual is the sellers' side of the market. It is not. Buyers and sellers both live in the primal. The dual is the price side, and the price belongs to the market rather than to either party.

4.4 Convex and non-convex markets, and why a supporting price exists here

Everything above depends on a property that is easy to take for granted: that there exists a single price at which every accepted order is willing to trade and every rejected order is not. That property is not automatic. It is a consequence of convexity, and it is worth being explicit about what convexity means here because the whole sequence of the paper turns on it.

A set is convex if, whenever two points belong to it, everything on the straight line between them belongs to it too. The feasible set of the programme above is defined by linear equalities, linear inequalities and box bounds, which makes it a polyhedron, and a polyhedron is convex: blend any two feasible allocations in any proportion and the blend is feasible. The objective is linear, which makes it simultaneously concave and convex, and a concave objective maximised over a convex set is the textbook case in which strong duality holds. The primal optimum and the dual optimum are equal, there is no gap between them, and the KKT conditions are not merely necessary but sufficient to characterise optimality.

The practical content of that abstract statement is exactly what the exchange needs. Strong duality is the guarantee that the multipliers π and μ exist at all, that they take finite values, and that the case table above is a complete description of the outcome rather than an approximation to it. The uniform price is not something the engine searches for and hopes to find. It is delivered by the same solve that produced the allocation, and its existence was assured before the solver was started. A market with this structure can be called convex, and in a convex market the question “is there a price that explains this outcome” never has to be asked, because the answer is always yes.

Divisibility is what buys this. Because the acceptance variable ranges over the whole interval from 0 to 1 rather than over the two isolated values 0 and 1, the feasible set has no holes in it, and the engine can always move a little in some direction and remain feasible. That is not only a theoretical nicety. If demand totals 13,000 MMBtu while the two cheapest sellers offer 11,000 between them and the next offers 9,000, then under all-or-nothing acceptance the achievable supply totals are 11,000 or 20,000.

Neither equals 13,000 and balance simply fails. The fraction is what allows supply to meet demand exactly rather than at the nearest achievable step, and it is simultaneously what allows an order to sit at the margin and thereby pin the price down.

Divisibility is also a market-design commitment rather than a modelling convenience, and Section 9.6 returns to it. A buyer who is partially filled has traded on the exchange. A buyer who was rejected outright has not, and may not return next session.

4.5 But in reality: when the market is not convex

The block order is the product Section 2.3 says may be missing, and it is also what makes the model hard. A block, or strip, is all-or-none across its days: take the whole profile, or nothing. Its acceptance variable therefore becomes a binary switch:

$$y \in \{0,1\}\ in\ place\ of\ x \in [0,1] \tag{11}$$

The set containing only 0 and 1 is not convex, so the feasible region acquires gaps. Strong duality fails, the case table of Section 4.3 no longer applies, and no supporting uniform price need exist. An accepted block may be out of the money at the published price, or a profitable block may have to be rejected.

This non-convexity is a fact about gas, not an artefact of the model. A single LNG cargo cannot be split. A take-or-pay minimum guaranteed quantity is a contractual floor. A plant with a continuous process needs gas on every day of a run or on none. Modelling the indivisibility away does not remove it; it only moves it to the telephone call the participant makes instead of submitting an order. A harder problem that clears the right trade is better than an easy problem nobody brings orders to.

This is why the paper is sequenced as it is. The linear programme establishes the objective, the acceptance rules, the price convention and the treatment of capacity where all of them are guaranteed to be well defined. Section 5 then sets out the main model, a mixed-integer linear programme solved by branch and cut, and measures what is given up against that benchmark.

So what does an exchange do when the market contains many such indivisibilities, so that no single price may explain the outcome? It keeps the linear programme as the foundation, and builds on top of it a mixed-integer linear programme, solved by branch and cut. That is the main model of this paper, and Section 5 sets it out.

5. The Main Model: MILP and Branch and Cut

Section 4 cleared a market in which every order is divisible, so a supporting price is guaranteed to exist. Real gas markets are not like that. This section introduces the product the trade data suggests is missing, the multi-day block order, and sets out what the exchange must do once the guarantee of a supporting price is gone. Nothing here changes IGX's existing products, which remain divisible and are fully handled by the linear programme. The binary variables appear only for block orders.

5.1 Why a MILP, and in what sense it is non-linear

Every equation in the model stays linear. What changes is that some variables can only take the values 0 or 1. That makes the problem non-convex: the feasible region has gaps, and the reasoning of Section 4.3, which relied on nudging a variable a little, stops working.

In effect, the model behaves non-linearly even though it is written linearly. A block's delivered gas equals its switch multiplied by its quantity, and the switch jumps from 0 to 1 with nothing in between. A mixed-integer linear programme captures this jump exactly, using only linear equations and integer variables.

There is a second reason for choosing this route. Prices and quantities could in principle be found together by writing the primal constraints, the dual constraints and the complementarity conditions into one problem. That problem contains products of variables, such as a price multiplied by an acceptance fraction, so it is genuinely non-linear and hard to solve reliably at scale. The MILP route avoids this. Integer decisions settle who trades. A separate linear programme, solved once those decisions are fixed, settles prices.

$$B = \varnothing \;\Rightarrow\; LP, B \neq \varnothing \;\Rightarrow\; MILP \tag{12}$$

Under IGX's existing products alone the problem is a linear programme. It solves in well under a second at IGX's scale, the answer is provably optimal, and all multipliers come out directly. Once block orders are introduced it becomes a MILP. Both cases are handled by the same engine.

5.2 What a block order is

A block order, or strip, covers an arbitrary set of delivery days with a chosen quantity on each day, accepted or rejected as a single package. One switch governs every day of the profile:

$$y_o \in \{0,1\}, q_{o,d} = y_o \cdot MaxQty_{o,d} \qquad \forall d \in D(o) \tag{13}$$

The quantity on each day is the requested quantity multiplied by a single switch shared across all days. There is no per-day fraction and no per-day switch. That is precisely what makes it a package rather than a bundle of separate orders. The days need not be consecutive or uniform; what matters is that the participant has declared the package to be a single commercial decision.

5.3 The full model

Sets. Delivery points n. Gas days d. Divisible orders O, each eligible at one or more points $N(o)$. Block orders B, each at one point $n(b)$ over a set of days $D(b)$.

Variables. Acceptance fractions $x_{o,n}$ between 0 and 1 for divisible orders at each eligible point. Switches $y_b \in \{0, 1\}$ for block orders. Slack variables $s^D_{n,d}$ and $s^S_{n,d}$ for unserved demand and unsold supply at the price bounds P^{min} and P^{max}.

Objective. Welfare, with $\sigma = +1$ for buy orders and -1 for sell orders:

$$max\ W = \sum_{o \in O} \sum_{n \in N(o)} \sigma_o\, p_{o,n}\, q_o\, x_{o,n} + \sum_{b \in B} \sigma_b\, p_b \sum_{d \in D(b)} q_{b,d}\, y_b$$

$$- \sum_{n,d} \left(P^{max} s^D_{n,d} - P^{min} s^S_{n,d} \right) \tag{14}$$

subject to

$$Sell_{n,d} - Buy_{n,d} + s^D_{n,d} - s^S_{n,d} = 0 \qquad \forall n, d \ [\pi_{n,d}] \tag{15}$$

$$Sell_{n,d} \le K_{n,d} \qquad \forall n, d \ [\mu_{n,d}] \tag{16}$$

$$\sum_{n \in N(o)} x_{o,n} \le 1 \qquad \forall o \in O \ [\lambda_o] \tag{17}$$

$$0 \le x_{o,n} \le 1, y_b \in \{0,1\} \qquad \forall o, n, b \tag{18}$$

In words: maximise welfare, so that at every point and day gas in equals gas out, injection stays within capacity, no multi-point order is filled more than once in total, and each block is either fully on or fully off. Here $Sell_{n,d}$ and $Buy_{n,d}$ collect the accepted quantities of all divisible and block orders at that point and day. Orders priced outside the band are rejected at entry. The penalty on the slack variables keeps every clearing price inside the band through the dual, and keeps the problem feasible for any order book. With no block orders, this is exactly the linear programme of Section 4, applied at many points in Section 10.

5.4 What the primal solves and what the dual solves

Table 5. What the primal solves and what the dual solves.

	Primal	**Dual**
Question	Who trades, and how much gas moves?	What is each constraint worth, and what prices support the outcome?
Variables	x and y (acceptance)	π (price), μ (capacity value), λ (order's best surplus)
In the LP	Solved together with the dual in one solve	Exists and is exact, by strong duality

	Primal	Dual
In the MILP	Solved by branch and cut	Not defined while y is free; found after y is fixed

Everything in Section 4.3 relies on being able to differentiate, on asking what happens if a variable is nudged slightly. A block order's feasible set is two isolated points, the whole profile and nothing. There is no slightly more than zero, and there is generally no multiplier vector that supports the integer optimum. A solver may still report numbers that look like duals, but they belong to one node of the search tree and have no claim to being prices.

The engine therefore never reads prices off a problem that still contains switches. It fixes the switches first:

$$solve\ MILP \ \rightarrow \ y^* \ \rightarrow \ fix\ y = y^* \ \rightarrow \ solve\ LP(y^*) \ \rightarrow \ \pi, \mu$$

Once every switch is locked, accepted blocks become ordinary fixed quantities and the remaining problem is a plain linear programme. The feasible set is convex again, the KKT conditions hold, strong duality holds, and the prices are meaningful. The prices obtained are the prices that support that particular combination of switches; a different combination would give different prices. In a convex market, allocation and price come out of one solve. In a non-convex market they come out of two, and the second is conditional on the first.

Section 5.12 onwards writes this whole system out in full: the primal, the dual, every complementarity condition, and a single mixed-integer programme that contains all of them.

5.5 The test each accepted block must pass

For divisible orders, the property that no accepted order is out of the money is automatic. For block orders it is not, and must be tested explicitly. For a block accepted with its switch on:

$$\sum_{d \in D(b)} \left(p_b - \pi_{n,d}\right) q_{b,d} \geq -\varepsilon \qquad buy\ block\ b \quad (19)$$

$$\sum_{d \in D(b)} \left(\tau_{n,d} - p_b\right) q_{b,d} \geq -\varepsilon \qquad sell\ block\ b \quad (20)$$

Add up, across every day of the block, the profit or loss the participant makes at the published prices. The total must not be negative, up to the small tolerance ε of Section 5.10.

Why this is judged over the whole profile. A block is a single commercial decision covering a period. It may be under water on individual days provided it comes out whole overall. That is what buying a package means. Testing day by day would reject blocks the participant is perfectly happy with, and would defeat the purpose of offering the product.

Why violation is forbidden. If the sum is negative, the participant has been made worse off by being accepted. Correcting that would require a side payment, which Section 7.4 commits against. IGX has no mechanism for side payments, and creating one would be a major market-design change requiring regulatory approval.

5.6 Branch and cut: the allocation-and-pricing loop

The full procedure is a loop, shown on the right of Figure 2.

- **Solve the MILP.** Branch and bound relaxes the switches to lie between 0 and 1, solves that linear programme, and branches on any switch that comes out fractional, one sub-problem with it at 0 and one with it at 1. This finds the welfare-best combination of accepted blocks.
- **Lock the switches** at the chosen combination y*, and solve the price LP that remains to read π and μ.
- **Select the published price** by the midpoint rule of Section 7.5.
- **Test every accepted block** as in Section 5.5. If all pass, publish. If any fails, add a no-good cut and solve again.

The cut forbids exactly the combination just rejected:

$$\sum_{o:\, y_o^*=1} (1 - y_o) + \sum_{o:\, y_o^*=0} y_o \geq 1 \tag{21}$$

The first sum counts switches that were on and are now off; the second counts switches that were off and are now on. Requiring the total to be at least 1 forbids repeating the identical combination. Because the cut is added inside the search, this is branch and cut: branching handles integrality, and cuts remove combinations that admit no consistent price. In the language of Section 1.4, it is a combinatorial Benders cut, with the MILP as the master problem and the price LP as the subproblem.

Why the loop always ends. There are finitely many block orders and therefore finitely many combinations of switches. Each iteration eliminates one combination permanently. In the worst case the engine reaches the combination that accepts no blocks at all, which passes trivially because it has no accepted blocks to test. Clearing no blocks is always feasible.

What the loop costs. Each re-solve adds a constraint, so welfare can only fall:

$$W_{published} \leq W_{best}, \Delta W = W_{best} - W_{published} \tag{22}$$

ΔW is the cost of uniform pricing: the welfare given up in exchange for one clean price per point and day instead of compensating individuals separately. It is reported every session. A persistently large ΔW is evidence that the block product is badly designed, and the evidence base for revisiting the no-side-payments commitment if it proves expensive. When no block orders are present, ΔW is zero by construction.

5.7 A worked example

One point, Dahej, two gas days. One block buyer wants 100 MMBtu on both days at Rs 1,000.

Table 6. Order book for the worked example: one delivery point, two gas days.

Order	Day	Side	Quantity (MMBtu)	Price (Rs/MMBtu)
Block buyer	1 and 2	Buy	100 each day	1,000
S1	1	Sell	100	500
S1b	1	Sell	50	700
D1	1	Buy	50	600
S2	2	Sell	100	900
D2	2	Buy	50	1,450

Round 1. The MILP accepts the block, for welfare of Rs 60,000. The block takes all of S1 on day 1 and all of S2 on day 2, so D1 and D2 receive nothing.

Price test. On day 1, D1 and S1b are rejected, so the price lies between Rs 600 and Rs 700. On day 2, D2 is rejected while bidding Rs 1,450, so the price must be at least Rs 1,450. The best the block can do is (1,000 − 600) × 100 + (1,000 − 1,450) × 100 = −Rs 5,000. It loses money at every valid price, so it fails.

Round 2. A no-good cut forbids accepting the block. The MILP now clears D1 against half of S1, and D2 against half of S2. Welfare is Rs 32,500. Prices are Rs 500 on day 1 and Rs 900 on day 2, set by the partially filled sellers.

Result. The block is rejected, although at these prices it would have earned Rs 60,000. It is reported as paradoxically rejected. The welfare given up for one clean price is ΔW = 60,000 − 32,500 = Rs 27,500, and it is published.

5.8 Uniform buyer price, multi-point sellers and opportunity cost

The market has a specific shape. Buyers at a point all pay one uniform price, $\pi_{n,d}$. Sellers who can deliver at several points face several prices, one per point, and the engine decides where they clear. This is where the division of labour between primal and dual, and the role of opportunity cost, becomes concrete.

Where a multi-point seller clears. Stationarity for seller o at point n gives:

$$q_o\left(\pi_{n,d} - \mu_{n,d} - p_{o,n}\right) - \lambda_o \leq 0 \qquad = 0 \; if \; x_{o,n} > 0 \quad (23)$$

The seller clears only at the point that gives it the highest surplus, and its dual λ_o equals that best surplus. λ_o is therefore the seller's opportunity cost: what it would give up by selling anywhere else. Multi-point buyers are treated identically.

Table 7. Dual variables and the opportunity costs they measure.

Dual or measure	What it is	Opportunity cost of
μ at a point and day	Value of one more unit of throughput	Scarce delivery capacity
λ for a multi-point order	Best surplus across eligible points	Trading at a worse point
ΔW per session	Welfare lost to keep one clean price	Uniform pricing with no side payments

The role of volume maximisation. Welfare is maximised first. Among all allocations with that welfare, the engine then picks the one that moves the most gas (Section 7.6). This second stage never lowers welfare. It reduces a different opportunity cost: capacity and willing trades left unused when several allocations tie. Its duals are discarded, and prices still come from the admissible range.

Taken together, the primal keeps volume as high as welfare allows, and the dual puts a rupee value on each source of lost opportunity: capacity, location and the pricing rule. Published together, they show what cleared, what the market could have gained, and why it did not.

5.9 Paradoxical rejection is permitted; branches are rejected, not participants

For a block whose switch is off, if

$$\sum_{d \in D(b)} \left(p_b - \pi_{n,d}\right) MaxQty_{b,d} > 0 \tag{24}$$

then the order would have been profitable at the published prices, yet was not accepted. This is a paradoxically rejected block. It is recorded and reported to the participant. It is not an error and it is not corrected.

Why the treatment is asymmetric. Symmetric treatment would insist that every profitable order be accepted and every accepted order be profitable. With all-or-nothing orders present there may exist no set of uniform prices under which every in-the-money block is accepted. This is a structural property of non-convex problems, and mature European power exchanges permit and report paradoxically rejected blocks for exactly this reason. If symmetry is enforced anyway, the engine rejects one candidate after another, searching for a solution that does not exist, and ends at a badly degraded allocation or at nothing traded. The rule is: forbid acceptance that harms, permit rejection that disappoints, and report the latter.

Branch rejection, not participant rejection. When an accepted block fails the test, the engine rejects the branch, meaning the combination of switches, and tries another. Every order remains in the market and may be accepted under a different combination. Participant rejection, removing a participant's orders so the calculation succeeds, answers a different question, what the price is among the remaining participants, and nobody is told. Its symptom is a progressive, unexplained collapse in participation across runs, with no error and no audit trail. That is why this distinction belongs in the specification rather than being left to implementation.

5.10 Numerical tolerances and the solve-time budget

Computers store numbers approximately. Without written rules, one implementation calls a borderline order accepted and another calls it rejected on the same order book. For blocks this matters more: a block profitable by a hundred-thousandth of a rupee could be judged unprofitable by numerical noise, triggering a branch rejection and changing the whole session.

The tolerance on the block surplus test is therefore deliberately negative: a block is treated as loss-making only when its profile surplus is below $-\varepsilon$. The asymmetry is intentional and belongs in the

specification rather than in solver settings. Rounding happens once, at publication. Any tiny imbalance is absorbed against the marginal order at that delivery point, never spread across participants.

An unfinished search must never be published as though it were optimal. If the solve-time budget is exceeded, the engine falls back to the best fully price-supported result found so far, recording the optimality gap. If none exists, the session is declared uncleared.

5.11 What the engine publishes

For each order: what was submitted; how much cleared on each day; accepted or rejected; and for block orders, whether the switch was on, the total surplus across the profile, and whether it was paradoxically rejected.

For each delivery point and day: the published price, the admissible range it was selected from, total volume cleared, and the capacity value where a limit applied.

For each session: published welfare, best possible welfare, the gap ΔW, the number of branch iterations, the list of paradoxically rejected blocks, the solver version and the input fingerprint.

The reporting unit is never an aggregate such as Monthly cleared 1.2 million. It is the order and its days: order M17, Monthly, Dahej, accepted, 40,000 MMBtu cleared on each of its 30 delivery days at the published Dahej price for each of those days.

5.12 The full primal-dual formulation: notation

Sections 5.3 to 5.6 describe the model and its solution loop in a compact form. This subsection and the next five write the IGX clearing problem out in full, in the style of the primal-dual framework of Madani and Van Vyve (2015b): the primal, the dual for a fixed block selection, every complementarity condition, and finally one mixed-integer programme that contains all of them. Writing everything out has two benefits. A referee can check, line by line, that the published prices satisfy every equilibrium condition. And a solver can work on one model and report a proven optimality gap, rather than on a loop.

To keep the system compact we use signed quantities, as Madani and Van Vyve do: a buy order has $Q > 0$ and a sell order has $Q < 0$. Then $p_o Q_o$ is the buyer's value or minus the seller's cost, and one formula serves both sides.

Table 8. Notation for the primal-dual formulation.

Symbol	Meaning
$n \in N,\ d \in D$	Delivery points and gas days
$i \in I$	Divisible single-point orders; i sits at point n(i) on day d(i)
$g \in G$	Multi-point orders; g is eligible at points N(g) on day d(g)
$b \in B$	Block orders at point n(b) over days D(b), with daily quantities $Q_{b,d}$
I^+, I^- (and $G^+, B^+, \ldots$)	Buy (Q > 0) and sell (Q < 0) subsets
$Q_i,\ Q_g,\ Q_{b,d}$	Signed quantities (MMBtu)

Symbol	Meaning
p_i, $p_{g,n}$, p_b	Limit prices (Rs/MMBtu); a multi-point order may state one price per point
$K^B_{n,d}$, $K^S_{n,d}$	Residual offtake and injection capacity
P^{min}, P^{max}	Price band
x_i, $x_{g,n} \in [0,1]$	Accepted fractions
$y_b \in \{0,1\}$	Block switches
$s^D_{n,d}$, $s^S_{n,d} \geq 0$	Unserved demand and unsold supply at the price bounds
$\pi_{n,d}$	Clearing price (dual of balance)
$\mu^B_{n,d}$, $\mu^S_{n,d} \geq 0$	Values of offtake and injection capacity
s_i, s_g, $s_b \geq 0$	Order surpluses (duals of the fill limits)
d^a_b, $d^r_b \geq 0$	Bounds on the loss of an accepted block and the opportunity cost of a rejected one

It is convenient to define, for any order o at point n on day d, the capacity-adjusted value of its quantity at market prices. A buyer pays the price plus the offtake value; a seller receives the price less the injection value:

$$\kappa_{o,n,d} = Q_o\,\pi_{n,d} + |Q_o|\,\mu^{\star}_{n,d}, \mu^{\star} = \mu^B \; if \; Q_o > 0, \;\; \mu^{\star} = \mu^S \; if \; Q_o < 0 \qquad \forall o,n,d \quad (25)$$

5.13 The primal programme

The primal maximises welfare over all orders, with the slack variables of the price band:

$$max\; W = \sum_{i \in I} p_i\,Q_i\,x_i + \sum_{g \in G} \sum_{n \in N(g)} p_{g,n}\,Q_g\,x_{g,n} + \sum_{b \in B} p_b \sum_{d \in D(b)} Q_{b,d}\,y_b - \sum_{n,d} \left(P^{max} s^D_{n,d} - P^{min} s^S_{n,d}\right) \quad (26)$$

subject to

$$x_i \leq 1 \qquad \forall i \in I\;[s_i] \quad (27)$$

$$\sum_{n \in N(g)} x_{g,n} \leq 1 \qquad \forall g \in G\;[s_g] \quad (28)$$

$$y_b \leq 1 \qquad \forall b \in B\;[s_b] \quad (29)$$

$$\sum_{i \in I_{n,d}} Q_i\,x_i + \sum_{g \in G_{n,d}} Q_g\,x_{g,n} + \sum_{b \in B_n} Q_{b,d}\,y_b - s^D_{n,d} + s^S_{n,d} = 0 \qquad \forall n,d\;[\pi_{n,d}] \quad (30)$$

$$\sum_{i \in I^{+}_{n,d}} Q_i\, x_i + \sum_{g \in G^{+}_{n,d}} Q_g\, x_{g,n} + \sum_{b \in B^{+}_{n}} Q_{b,d}\, y_b \le K^{B}_{n,d} \qquad \forall n, d \; [\mu^{B}_{n,d}] \quad (31)$$

$$\sum_{i \in I^{-}_{n,d}} |Q_i|\, x_i + \sum_{g \in G^{-}_{n,d}} |Q_g|\, x_{g,n} + \sum_{b \in B^{-}_{n}} |Q_{b,d}|\, y_b \le K^{S}_{n,d} \qquad \forall n, d \; [\mu^{S}_{n,d}] \quad (32)$$

$$x_i,\; x_{g,n},\; s^{D}_{n,d},\; s^{S}_{n,d} \ge 0 \quad (33)$$

$$y_b \in \{0,1\} \qquad \forall b \in B \quad (34)$$

Equations (30) to (32) are the only places where orders interact. Because no constraint links one delivery point to another, the whole system separates by point, as Section 10.5 uses.

5.14 The dual programme for a fixed block selection

Take any split of the block orders into accepted and rejected sets, $B = B^{a} \cup B^{r}$, and fix the switches with two families of constraints:

$$-y_b \le -1 \qquad \forall b \in B^{a} \; [d^{a}_{b}] \quad (35)$$

$$y_b \le 0 \qquad \forall b \in B^{r} \; [d^{r}_{b}] \quad (36)$$

With the switches fixed, integrality (34) can be dropped and the problem is a linear programme. Its dual, written with the capacity-adjusted values (25), is:

$$\min \sum_{i} s_i + \sum_{g} s_g + \sum_{b} s_b + \sum_{n,d} \left(K^{B}_{n,d}\, \mu^{B}_{n,d} + K^{S}_{n,d}\, \mu^{S}_{n,d}\right) - \sum_{b \in B^{a}} d^{a}_{b} \quad (37)$$

subject to

$$s_i + \kappa_{i,n(i),d(i)} \ge p_i\, Q_i \qquad \forall i \; [x_i] \quad (38)$$

$$s_g + \kappa_{g,n,d(g)} \ge p_{g,n}\, Q_g \qquad \forall g,\; n \in N(g) \; [x_{g,n}] \quad (39)$$

$$s_b + d^{r}_{b} + \sum_{d \in D(b)} \kappa_{b,n(b),d} \ge p_b \sum_{d \in D(b)} Q_{b,d} \qquad \forall b \in B^{r} \; [y_b] \quad (40)$$

$$s_b - d^{a}_{b} + \sum_{d \in D(b)} \kappa_{b,n(b),d} \ge p_b \sum_{d \in D(b)} Q_{b,d} \qquad \forall b \in B^{a} \; [y_b] \quad (41)$$

$$P^{min} \le \pi_{n,d} \le P^{max} \qquad \forall n, d \; [s^{S}, s^{D}] \quad (42)$$

$$s_i,\; s_g,\; s_b,\; \mu^{B}_{n,d},\; \mu^{S}_{n,d},\; d^{a}_{b},\; d^{r}_{b} \ge 0, \qquad \pi_{n,d}\; free \quad (43)$$

Each dual constraint belongs to one primal variable, shown in brackets. Constraint (42) is the price band, and it appears here, rather than as a rule, because of the penalty on the slacks in (26): the dual constraint of s^{D} is $\pi \le P^{max}$ and that of s^{S} is $\pi \ge P^{min}$.

5.15 Complementarity: the equilibrium conditions

For a fixed block selection, a primal point and a dual point are both optimal if and only if they satisfy the following complementarity conditions. Each says that a dual variable and the slack of its constraint cannot both be positive.

$$s_i\,(1 - x_i) = 0 \qquad \forall i \quad (44)$$

$$s_g\left(1 - \sum_{n \in N(g)} x_{g,n}\right) = 0 \qquad \forall g \quad (45)$$

$$s_b\,(1 - y_b) = 0 \qquad \forall b \quad (46)$$

$$\mu^B_{n,d}\left(K^B_{n,d} - Buy_{n,d}\right) = 0 \qquad \forall n, d \quad (47)$$

$$\mu^S_{n,d}\left(K^S_{n,d} - Sell_{n,d}\right) = 0 \qquad \forall n, d \quad (48)$$

$$x_i\left(s_i + \kappa_{i,n(i),d(i)} - p_i\,Q_i\right) = 0 \qquad \forall i \quad (49)$$

$$x_{g,n}\left(s_g + \kappa_{g,n,d(g)} - p_{g,n}\,Q_g\right) = 0 \qquad \forall g,\ n \in N(g) \quad (50)$$

$$y_b\left(s_b - d^a_b + \sum_d \kappa_{b,n(b),d} - p_b \sum_d Q_{b,d}\right) = 0 \qquad \forall b \in B^a \quad (51)$$

$$y_b\,d^r_b = 0, \qquad (1 - y_b)\,d^a_b = 0 \qquad \forall b \quad (52)$$

$$s^D_{n,d}\left(P^{max} - \pi_{n,d}\right) = 0, \qquad s^S_{n,d}\left(\pi_{n,d} - P^{min}\right) = 0 \qquad \forall n, d \quad (53)$$

Each condition has a plain reading. Conditions (44) and (49) together are the case table of Section 4.3: an order with positive surplus is fully filled, an order that is filled at all is at or in the money, and so only an order exactly at the money can be partly filled. Conditions (47) and (48) say capacity has value only when it is full. Condition (50) says a multi-point order clears only at points where it earns its best surplus. Condition (53) says curtailment happens only at the edge of the price band.

Lemma 1 (meaning of d^a and d^r). Let a primal and a dual point satisfy (27) to (36), (38) to (43) and (44) to (53). Define the profile surplus of block b at the given prices as

$$\Sigma_b = \sum_{d \in D(b)} \left(p_b\,Q_{b,d} - \kappa_{b,n(b),d}\right) \qquad \forall b \quad (54)$$

Then (i) for an accepted block, d^a_b bounds its loss, and (ii) for a rejected block, d^r_b bounds its opportunity cost:

$$d^a_b \geq max[0, -\Sigma_b] \qquad \forall b \in B^a \quad (55)$$

$$d^r_b \geq max[0, \Sigma_b] \qquad \forall b \in B^r \quad (56)$$

Proof. (i) For an accepted block, (51) gives $s_b - d^a_b = \Sigma_b$. Since $s_b \geq 0$, any loss is at most d^a_b. (ii) For a rejected block, (46) gives $s_b = 0$, and (40) then reads $d^r_b \geq \Sigma_b$, with $d^r_b \geq 0$. ■

Lemma 2 (multi-point orders). For a multi-point order, the surplus equals its best surplus across its eligible points, and it clears only at a point that attains it:

$$s_g = max\left[0,\ \max_{n \in N(g)} \left(p_{g,n}\, Q_g - \kappa_{g,n,d(g)}\right)\right] \qquad \forall g \quad (57)$$

Proof. (39) gives the inequality for every eligible point; (45) and (50) make it an equality at any point where the order clears, and force $s_g=0$ if it does not clear at all. ■ This is the opportunity cost of Section 5.8 written as an equation.

5.16 One MILP for the whole market: the IGX feasible set

The difficulty is that the best split $B^a \cup B^r$ is not known in advance. The device of Madani and Van Vyve (2015b) removes it. Weak duality says the primal objective (26) can never exceed the dual objective (37). So if we add one constraint saying it is at least as large, the two are forced to be equal, and by strong duality that equality holds exactly when all of (44) to (53) hold. None of the complementarity conditions, which are products of variables, has to be written into the model.

The IGX feasible set (IGX-FS) consists of the primal constraints (27) to (34), the dual constraints (38), (39), (42) and (43), and the following:

$$W \geq \sum_i s_i + \sum_g s_g + \sum_b s_b + \sum_{n,d} \left(K^B_{n,d}\, \mu^B_{n,d} + K^S_{n,d}\, \mu^S_{n,d}\right) - \sum_{b \in B} d^a_b \qquad (58)$$

$$s_b + d^r_b - d^a_b + \sum_{d \in D(b)} \kappa_{b,n(b),d} \geq p_b \sum_{d \in D(b)} Q_{b,d} \qquad \forall b \in B\ [y_b] \quad (59)$$

$$d^r_b \leq M_b\,(1 - y_b) \qquad \forall b \in B \quad (60)$$

$$d^a_b \leq M_b\, y_b \qquad \forall b \in B \quad (61)$$

$$d^a_b = 0 \qquad \forall b \in B \quad (62)$$

Constraint (58) is the strong-duality constraint. Constraint (59) merges (40) and (41) into one family, and the big-M constraints (60) and (61) dispatch each block to the right one: when $y_b=1$ only d^a_b may be positive, and when $y_b=0$ only d^r_b. Constraint (62) is the IGX market rule of Section 7.4: no block may be paradoxically accepted. The constants must be large enough not to cut off any valid price. Since every price lies in the band and, by (38), every capacity value is at most P^{max}-P^{min}, it is sufficient to take

$$M_b = 2\left(P^{max} - P^{min}\right) \sum_{d \in D(b)} \left|Q_{b,d}\right| \qquad \forall b \in B \quad (63)$$

Proposition 1 (equivalence). (I) If a point is feasible for IGX-FS, then setting $B^a=\{b:y_b=1\}$ and $B^r=\{b:y_b=0\}$, its primal and dual parts satisfy every condition (27) to (53): the prices support the allocation, every divisible and multi-point order is at equilibrium, and no accepted block loses money. (II) Conversely, any allocation and prices that satisfy those conditions for some block split, with prices inside the band, can be extended with values of d^a and d^r to a feasible point of IGX-FS.

Proof sketch. (I) The big-M constraints split (59) into (40) and (41), so the dual is feasible for the fixed split, and (58) together with weak duality gives equality of (26) and (37), which by strong duality is equivalent to (44) to (53); (62) and Lemma 1 give no loss. (II) Strong duality gives the equality of objectives; set $d^r_b = max[0, \Sigma_b]$ for rejected blocks and all other d to zero, which satisfies (60) and (61) by the choice (63). ■ This mirrors Theorem 1 of Madani and Van Vyve (2015b).

Using (62), the variables d^a and d^r can be eliminated altogether, exactly as Madani and Van Vyve do for PCR-FS. The block constraint and the strong-duality constraint become:

$$s_b + M_b\,(1 - y_b) + \sum_{d \in D(b)} \kappa_{b,n(b),d} \geq p_b \sum_{d \in D(b)} Q_{b,d} \qquad \forall b \in B\ [y_b] \quad (64)$$

$$W \geq \sum_{i} s_i + \sum_{g} s_g + \sum_{b} s_b + \sum_{n,d} \left(K^B_{n,d}\,\mu^B_{n,d} + K^S_{n,d}\,\mu^S_{n,d}\right) \quad (65)$$

The clearing problem is then a single MILP with no auxiliary variables and no more binary variables than there are block orders:

$$\max_{IGX-FS} W \quad (66)$$

The branch-and-cut loop of Section 5.6 is one way to solve it. The MILP of Section 5.3 plays the role of a master problem over the switches, the price LP plays the role of the subproblem, and each no-good cut removes one switch combination that cannot be supported by prices. Solving (66) directly lets an off-the-shelf solver do the same search inside one tree and report a proven bound.

5.17 Other objectives on the same feasible set

Because IGX-FS describes every valid market outcome, any objective can be optimised over it. Four are useful for IGX.

Welfare first, volume second. Let W^* be the optimal value of (66). The second stage of Section 7.6 is:

$$max \sum_{i \in I^+} Q_i\,x_i + \sum_{g \in G^+} \sum_{n} Q_g\,x_{g,n} + \sum_{b \in B^+} \sum_{d} Q_{b,d}\,y_b$$

$$over\ IGX - FS, subject\ to\ W \geq W^* - \delta \quad (67)$$

Here δ is a small tolerance. Madani and Van Vyve show that, once blocks are present, the volume-maximising and welfare-maximising outcomes can differ; (67) keeps welfare optimal and then moves the most gas.

Minimum opportunity cost. Keeping the version of IGX-FS with d^r, the outcome that minimises the total opportunity cost of paradoxically rejected blocks is:

$$\min_{IGX-FS} \sum_{b \in B} d^r_b \quad (68)$$

The admissible price range. With the switches fixed at their optimal values $\hat{y}$ and welfare at its optimum $\hat{W}$, the lowest and highest valid prices at each point and day solve two linear programmes, and the published price is their midpoint (Section 7.5):

$$\underline{\pi}_{n,d} = min\{\pi_{n,d}:\ IGX - FS,\ y = \hat{y},\ W = \widehat{W}\} \qquad \forall n,d \quad (69)$$

$$\underline{\pi}_{n,d} = max\{\pi_{n,d}:\ IGX - FS,\ y = \hat{y},\ W = \widehat{W}\} \qquad \forall n,d \quad (70)$$

$$\pi_{n,d}^{pub} = \frac{1}{2}\left(\underline{\pi}_{n,d} + \underline{\pi}_{n,d}\right) \qquad \forall n,d \quad (71)$$

Where the welfare goes. At any feasible point of IGX-FS, (58) holds with equality. Read as an identity, it splits welfare into the surplus of every participant, the rent earned by scarce capacity, and any compensation that would be owed to paradoxically accepted blocks. The three bracketed terms are, in order, participants' surplus, capacity rent and PAB compensation:

$$W = \left(\sum_{i} s_i + \sum_{g} s_g + \sum_{b} s_b\right) + \left(\sum_{n,d} K_{n,d}^{B}\mu_{n,d}^{B} + K_{n,d}^{S}\mu_{n,d}^{S}\right) - \sum_{b} d_b^a \qquad (72)$$

Under IGX rules the last term is zero by (62). Relaxing (62) can only enlarge the feasible set, so welfare with compensated paradoxical acceptance is never lower, the result first noted by O'Neill, Sotkiewicz and Rothkopf (2007):

$$\max_{IGX-FS} W \ \leq \ \max_{IGX-FS\ without\ d^a=0} W \qquad (73)$$

Finally, the cost of uniform pricing published each session is the gap between the best welfare the primal alone could reach and the welfare of the published outcome:

$$\Delta W = \max_{(P1)-(P8)} W \ - \ \max_{IGX-FS} W \ \geq \ 0 \qquad (74)$$

When there are no block orders, the primal is an LP, strong duality holds without (58) needing to bind any switch, and $\Delta W=0$. Every quantity in Section 5.11 that the engine publishes is a variable or an objective value of this one system.

6. Key Concepts

Each term below is defined in the language of a gas exchange rather than abstractly. Sections 6.4 to 6.8 restate at greater length what Sections 4.2 to 4.5 introduced, and a reader who has absorbed the introduction can skim them. The order follows the build sequence of the reference implementation, which constructs the problem one piece at a time on a market small enough to verify by hand before the same machinery is pointed at a full order book.

6.1 The double-sided auction

Both sides submit orders. A buyer states a limit price, the most they will pay per MMBtu, and a requested volume. A seller states a floor price, the least they will accept, and an available volume. In an auction all orders for a delivery period are collected and cleared at one moment rather than matched as they arrive.

An IGX order carries at minimum the delivery point, contract type, volume, limit price, and margin sufficient to back the notional. The only quantity the exchange chooses is how much of each order to accept. Continuous matching under price and time priority, which IGX runs for intraday contracts, produces a sequence of bilateral trades at different prices rather than one clearing price and is outside the scope of this paper.

6.2 Uniform pricing

Every accepted order settles at one price, whatever it bid. The alternative, pay-as-bid, charges each buyer their own number and so teaches them to write a number below their true valuation, after which the order book stops carrying information about what gas is worth. Under uniform pricing the price a participant pays is set by other participants' orders rather than their own, so bidding honestly cannot hurt them. Uniform pricing also makes settlement arithmetic simple and gives one publishable number per delivery point per period.

6.3 Divisibility and the acceptance fraction

For each order the engine chooses an accepted fraction between 0 and 1. A value of 1 is a full fill, 0 is a rejection, and 0.22 means 22 per cent filled.

The fraction is not a volume. The gas that moves is the requested quantity multiplied by the fraction, which can look different for everyone. This is the first place a reader misreads the model, and the implementation prints both columns side by side for that reason.

Fractions are usually necessary for feasibility rather than being a convenience. If demand totals 13,000 MMBtu and the two cheapest sellers offer 11,000 between them with the next offering 9,000, then under all-or-nothing acceptance the achievable supply totals are 11,000 or 20,000. Neither equals 13,000 and balance fails. The fraction is what allows supply to meet demand exactly rather than at the nearest achievable step.

6.4 Welfare, and why the price cancels

A buyer with limit price p buying at π gains the difference per MMBtu. Summed over accepted volume this is consumer surplus. A seller with floor p selling at π gains the difference the other way, and summed this is producer surplus. Welfare is their sum.

Expanding that expression collects the π terms into a bracket that is exactly the balance constraint, which is zero at any feasible allocation. Welfare therefore reduces to buyer value minus seller cost, with no price in it at all. The derivation is written out in Section 4.2.

This is the π-cancellation, and it is the structural fact the whole approach rests on. Money passing from buyer to seller is a transfer rather than a gain. It determines who captures the surplus, not how large the surplus is.

The consequence surprises people: the welfare-maximising allocation can be computed without knowing the clearing price, and the acceptance fractions fall out of that computation. Section 9.3 demonstrates it by recomputing the surplus split at three prices across the valid range. The split moves and the total does not.

Geometrically, consumer surplus is the area between the descending demand staircase and the price line out to the cleared volume, and producer surplus the area between the price line and the ascending supply staircase.

Figure placeholder. Demand and supply staircases with the price line, consumer surplus shaded above and producer surplus below, out to the cleared volume.

6.5 Primal and dual

Every constrained optimisation has a dual attached to it. The primal asks a physical question: who trades, and how much gas moves? The dual asks an economic one: what is each constraint worth, and what prices support that physical answer?

There is one dual variable per constraint. The dual of the balance constraint is the clearing price. The dual of the capacity constraint is the value of one more unit of capacity.

A common misreading is that the dual is the sellers' side of the market. It is not. Buyers and sellers both live in the primal. The dual is the price side, and the price belongs to the market rather than to either party.

6.6 KKT conditions

The Karush-Kuhn-Tucker conditions describe what an optimal solution looks like. They are primal feasibility, dual feasibility meaning multipliers of the right sign, complementary slackness meaning a non-binding constraint has zero shadow price, and stationarity meaning that no small change to any quantity would improve the total.

The last two together give the acceptance rules. A buyer bidding above the price is fully accepted, one bidding below is rejected, and one bidding exactly at the price is partially accepted. Sellers work the same way, measured against their floor plus the capacity dual.

For a linear programme these conditions are both necessary and sufficient for a global optimum. A set of values satisfying all four is provably optimal, and no better allocation exists anywhere in the feasible set.

The implementation checks all four numerically after each solve and prints pass or fail. This is a regression test on the software, not a test on the market. The solver stopped because these conditions hold, so a failure means the model or the data is wrong, and the session must halt.

6.7 Shadow prices

A shadow price answers how much the objective would improve if a constraint were relaxed by one unit. It is a rate of change. The balance dual is the price.

Suppose one more MMBtu appeared at the delivery point, free. It would go to the highest-valued unserved buyer and be worth that buyer's bid, and that buyer sits exactly at the market's margin. The value of one more unit and the clearing price are therefore the same number, reached from two directions. The implementation shows this empirically rather than asserting it: re-solving with the balance right-hand side nudged by various amounts, the welfare change per unit is constant and equals the reported dual.

The capacity dual measures scarce delivery room. If the constraint is slack it is zero and capacity is free. If it binds, it is the rupee value per day of one more MMBtu of throughput.

The capacity dual multiplied by the volume shortfall reproduces the welfare loss exactly. That identity closing numerically is what it means for a number to be a shadow price, and the implementation asserts it.

6.8 Convexity, and what it guarantees

A set is convex if everything on the line between any two of its members is also a member. The feasible set of the divisible problem is a polyhedron and therefore convex, and the objective is linear and therefore concave. Together these give strong duality: the primal and dual optima coincide, and a price system supporting the allocation is guaranteed to exist.

Non-convexity is what all-or-none acceptance introduces. A block order's feasible set is two isolated points, and there is nothing between them. Differentiation becomes meaningless, strong duality fails, and a supporting price may not exist at all. Section 5 is about what an exchange does in that situation.

The terms are worth keeping straight because they are sometimes used loosely. Concavity is a property of the objective and is what makes a maximum well-behaved. Convexity is a property of the feasible set and is what makes the multipliers exist. This market is convex in its divisible form, and it is convexity, not any property of gas itself, that is lost when block orders are admitted.

7. Design Decisions

Eight design decisions shape the engine. Each one sits at a single step of the clearing flow, so Figure 2 is the simplest way to see how they fit together. The number in brackets is the subsection where the decision is argued.

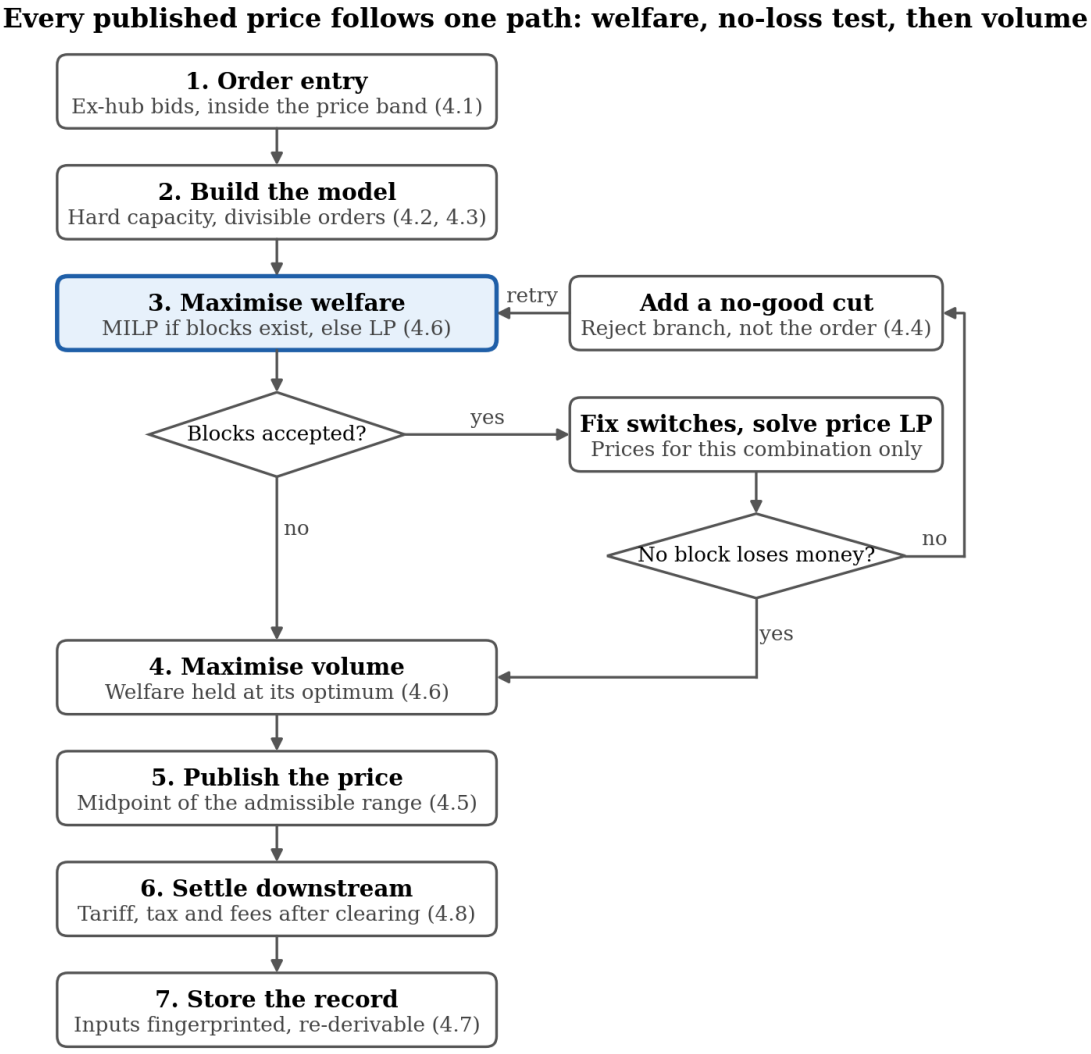


Figure 2. The clearing flow.
The pricing loop on the right runs only when block orders are accepted.

Read from top to bottom: orders enter as ex-hub bids and the model maximises welfare. Only if blocks are accepted does the pricing loop run, and a failed test removes a combination of blocks, never a participant. Volume is maximised only once welfare is fixed, and settlement happens after the price is published. The table summarises each decision; the subsections that follow give the full reasoning, so a later reader can see what reversing any one of them would cost.

Table 9. Summary of the design decisions.

#	Decision	In one line	Cost of reversing it
4.1	Bids are ex-hub	Participants net their own tariff and tax off their bid	Costs counted twice; no single uniform price
4.2	Capacity is a hard constraint	The pipe limit shapes the allocation, not just checks it	Trades cleared that cannot be delivered
4.3	Divisible orders by default	Partial fills keep the model an LP, so a price always exists	Supply and demand may fail to balance
4.4	One price, no side payments	No accepted order loses money; profitable rejections are reported	Needs an uplift fund and hidden charges
4.5	Midpoint price rule	One order book gives one price, whatever the solver	Two solvers could publish two prices
4.6	Welfare first, volume second	Most value first, then most gas among equal-value outcomes	Back to pure volume maximisation
4.7	Reproducibility	Same inputs give the same outputs, every time	Disputed prices cannot be re-derived
4.8	Settlement is downstream	Tariff, tax and surplus reports come after clearing	Clearing becomes hard to audit

Each is recorded with its reasoning so a later reader can see what reversing it would cost.

7.1 Bids are ex-hub

Limit prices are for gas at the delivery point, before transmission tariff, tax, marketing margin and fees. A buyer bidding ex-hub has already netted their downstream costs off what they are willing to pay, so counting those costs again inside the objective would double-count them and distort the allocation. They are applied after clearing instead. Reversing this makes the objective participant-specific, and a single uniform price becomes impossible across buyers with different offtake points.

7.2 Capacity is a hard constraint

Physical delivery limits go inside the optimisation rather than being checked afterwards. A limit applied after the fact can invalidate an allocation but cannot shape it, so the engine would produce an outcome it then has to partly withdraw. Where offered volume far exceeds throughput, capacity rather than the bid and ask crossing is what rations the market, and a price set without reference to it is the wrong price. Where a delivery point has no binding limit, the constraint is omitted entirely rather than set to a very large number, because a constraint that can never bind still contributes a spurious zero multiplier and enlarges the set of valid prices for no benefit.

7.3 Divisible orders by default

Every order that can be partially filled is, which keeps the problem a linear programme and guarantees that a supporting price exists. All-or-none acceptance is reserved for block orders, which is where its cost is justified. Participants must be told partial execution is possible, since an order for 23,000 MMBtu may fill to 13,500, and that belongs in the order-entry interface rather than only in technical documentation.

7.4 One price, no side payments

European and North American designs were both reviewed, and the European convention is adopted. An order that would lose money at the published price is never accepted. An order that would have been profitable may be rejected, and this is disclosed to the participant. The alternative is to accept loss-making orders and compensate them separately, which requires an uplift pool, a rule for who funds it, and a surcharge invisible at the moment of trading, so the published price stops being what anyone actually pays. At IGX's stage, confidence in the neutrality of clearing is the scarce resource, and no side-payment mechanism exists to build on. This costs nothing while all orders are divisible, and it is settled now because it determines the shape of the block-order extension in Section 5.

7.5 Price convention: the midpoint of the admissible range

In a real order book there is rarely one correct price. If the highest unmatched bid is Rs 1,000 and the lowest unmatched ask is Rs 1,020, any price in between clears exactly the same quantities and satisfies every KKT condition equally well. The mathematics does not choose between them. The solver does, arbitrarily, based on its internal method, and two solvers or two builds of the same solver can publish different prices for identical order books. For an exchange that is disqualifying.

The engine therefore computes the two ends of the admissible range for each point and day and publishes the midpoint, rounded to the market's price tick. Usually the range collapses to a single point, because a partially filled order pins the price to its own limit. Where it is genuinely wide, which happens only when no order is partially filled, the midpoint is published.

Four reasons. It is unique, so one order book yields one price independent of solver, machine or loading order. It is neutral, since publishing the lower end would systematically favour buyers and the upper end sellers. It is stateless, referring to no previous session and no external index, so it cannot be manipulated across sessions. And it is simple enough to explain to the participant disputing it.

It is worth being honest about what this rule is. The midpoint is not derived as uniquely correct. It is chosen, because it is neutral, unique and explainable. The rule belongs in IGX's published market rules, not only in a technical specification.

7.6 Welfare first, volume second

The objective is welfare, and among all allocations achieving maximum welfare the engine selects the one clearing the greatest volume. This is a lexicographic ordering, implemented as the two-stage solve of Section 9.2: the first stage maximises welfare, the second re-solves for maximum cleared volume with welfare pinned at its optimum by a constraint, and the duals of the second stage are discarded because its objective is no longer welfare.

Three reasons. Linear programmes frequently have multiple optimal solutions, especially where orders share limit prices, and an engine that leaves the choice among them to the solver's pivoting order is neither reproducible nor defensible to a participant. Gas is a quantity-driven market, and an exchange indifferent between clearing 90,000 and 91,850 MMBtu at equal welfare has misread what it is for. And the secondary criterion costs nothing in welfare terms by construction, since it operates only across allocations that are already welfare-optimal.

The objective itself is never adjusted to break the tie, because that would change what the market maximises. Remaining ties after both stages are broken by earliest order timestamp, then by order identifier, both stated in the rulebook. Reversing the ordering, volume first and welfare second, returns the engine to the objective this paper argues against.

7.7 Reproducibility is a market-rule obligation

The same order book, the same configuration and the same prior commitments must produce identical output every time. This requires the tie-break rules above, the price rule of Section 7.5, a solver version pinned and recorded per market release, a fixed order-loading sequence, and a stored fingerprint of each session's input. It is not an engineering nicety. A price disputed months later must be exactly re-derivable, and without that the exchange cannot demonstrate that the published price was the correct one.

7.8 Settlement is downstream

Consumer surplus, producer surplus, congestion rent allocation and hub-level aggregates are computed after clearing, from the published quantities and prices. They never enter the optimisation. The engine emits an unchangeable record and downstream systems consume it without calling back into clearing. Mixing settlement reporting into the allocation is the standard route to a clearing engine nobody can audit.

8. Preliminary Work: Mapping the Cost Stack

Before any clearing model was written, we built out what a molecule of gas actually costs between its origin and the burner tip.

A component catalogue records the individual cost elements: APM and HPHT domestic ceilings, long-term and spot LNG contract bases, basic customs duty by physical form, central excise, port wharfage and terminal handling at each of the eight operational LNG terminals, pipeline tariffs, and state-wise VAT. Each carries its payer, whether it varies regionally, its source and a confidence rating.

Two findings from that exercise shaped everything after it.

First, the exchange price is a minority of what the buyer pays. By the time gas reaches an industrial consumer, transmission tariff, marketing margin and tax have accumulated on top of it, and none of that is visible in the published clearing price.

Second, the tax treatment does not behave like a tax on a commodity. Natural gas sits outside GST and is taxed under state VAT at rates ranging from 5 per cent to roughly 14.5 per cent depending on the state and the offtake segment, while its transportation is a service inside GST at 18 per cent. The two cannot be set against one another, so the buyer absorbs both.

That is what motivated the decision to clear ex-hub and treat settlement as a separate downstream pass. The clearing engine should price the thing the market actually forms, and the cost stack sitting on top of it is a different object, formed by different processes, worth measuring separately rather than folding into the allocation.

9. The LP in Practice: A Single Delivery Point

This section runs the linear programme of Section 4 on order books at one delivery point. It confirms the mechanism against a market small enough to check by hand, then applies it to a realistic order book. What it reveals about capacity is what motivates multi-point orders in Section 10 and the market diagnosis in Section 11.

9.1 Idea

The base model asks the narrowest useful question. At one delivery point, for one delivery day, with every order divisible, what allocation maximises gains from trade, and what price supports it?

The point of starting here is that everything downstream depends on the price being derived rather than chosen, and that claim is only checkable on a market small enough to verify by hand. This section establishes the mechanism, confirms it against hand calculation, then applies it to a realistic order book and observes what happens. What it revealed was not what we expected, and it is what drove the rest of the paper.

9.2 Method

Setting. One delivery point, with Dahej as the running case. Arbitrary buy and sell orders for one delivery period. Every order divisible. Throughput capped at K, the available capacity for the period, operationally 75 per cent of nameplate.

Notation.

Table 10. Notation for the single-delivery-point programme.

Symbol	Meaning	Source
B, S	sets of buy and sell orders	order book
p, q (buy)	limit price and requested volume	order book
p, q (sell)	floor price and available volume	order book
K	available capacity for the period	pipeline operator
x	accepted fraction, 0 to 1	solver
π	clearing price per MMBtu	dual of balance
μ	capacity shadow price per MMBtu	dual of capacity
τ	effective seller threshold, $\pi - \mu$	derived
W	welfare for the period	objective value

The programme. Maximise welfare, defined as the sum of buyer value less the sum of seller cost, subject to gas bought equalling gas sold, total throughput at or below K, and every fraction between 0 and 1.

Three points on the constraints. Balance is an equality because there is no storage at the delivery point, so surplus supply would be gas in the pipe nobody ordered and surplus demand would be gas the exchange has promised and does not have. The lower bound on the fraction is not decorative, because without it the solver sets a negative fraction on an expensive seller, since subtracting a negative cost inflates welfare. Capacity is written on the sell side because sellers are the ones putting gas into the point, which makes the interpretation of μ direct.

Why earlier sessions do not appear. IGX clears sequentially, and gas for a given delivery day may already have been committed in earlier Monthly or Weekly sessions. That committed gas does not enter this session's balance, because a concluded trade is a matched pair: equal and opposite quantities at the same point and day, summing to zero. This holds because IGX contracts are compulsory physical delivery, non-transferable, day-specific, with physical imbalance handled by the transporter rather than the exchange. If any of those three were to change, a committed-position term would be needed in the balance.

Where the price comes from. Relaxing each constraint into the objective with its own multiplier and differentiating shows that the derivative of welfare with respect to accepting an order is that order's surplus. Combined with complementary slackness this yields the case table of Section 4.3.

The staged algorithm. Stage 1 solves the welfare programme and records the optimal welfare and the capacity dual. Stage 2 re-solves maximising cleared volume with welfare pinned by a constraint, records the allocation, and discards that stage's duals since the objective is now volume. Stage 3 constructs the admissible price interval from the allocation and the capacity dual, taking as the lower end the highest of all rejected or partial buyer limits and all accepted seller thresholds, and as the upper end the lowest of all accepted buyer limits and all rejected or partial seller thresholds, then publishes the midpoint. If the lower end exceeds the upper, the allocation is not optimal and the engine halts. Stage 4 applies transport, tax and fees downstream to convert the published price into each participant's delivered price.

Implementation. The reference implementation builds the model one piece at a time and then reads the allocation, price and shadow price off the same solve. Part 1 uses a hand-checkable market of three buyers and three sellers. Part 2 points the identical machinery at a synthetic order book of 15 sellers and 40 buyers at Dahej, with the point's capacity read from its pipeline sheet.

9.3 Results on a hand-checkable market

Buyers bid 1,000, 800 and 500 rupees for 100 MMBtu each. Sellers ask 300, 600 and 900 rupees for 100 MMBtu each.

Unconstrained. The optimum accepts the top two on each side, clearing 200 MMBtu for welfare of 90,000 rupees. No order is fractionally filled, so no order is marginal and the price is genuinely indeterminate across 600 to 800. The midpoint convention publishes 700, splitting the surplus 40,000 to buyers and 50,000 to sellers.

Recomputing the split at 600, 700 and 800 gives 60,000 and 30,000, then 40,000 and 50,000, then 20,000 and 70,000. The total is 90,000 every time. The price moves surplus between the sides without changing its size, which makes the choice of convention a distributional question rather than an efficiency one. This is the π-cancellation of Section 4.2, observed numerically.

With capacity of 150. Three things change together, none of them coded.

Table 11. The hand-checkable market without and with a capacity limit of 150.

	Unconstrained	**Capacity 150**
Welfare	90,000	80,000
Volume cleared	200	150
Clearing price π	700 (midpoint of 600 to 800)	800 (unique)
Capacity dual μ	0	200
Consumer surplus	40,000	20,000
Producer surplus	50,000	60,000

Welfare falls because the constraint costs real value. The price becomes unique because one buyer is now half filled, and a fractional fill pins the price to that order's limit. The capacity dual turns positive because capacity is now worth something. Scarcity raised the price and shifted surplus toward sellers.

The shadow price closes numerically: 200 multiplied by the shortfall of 50 is 10,000, exactly the welfare lost. That identity is the cleanest single check that the dual has been read correctly.

9.4 Results on the synthetic Dahej book

The same machinery, unchanged, applied to a book of 55 orders and two constraints.

Table 12. Inputs of the synthetic Dahej book.

Quantity	**Value**
Available capacity	91,850 MMBtu/day
Total sell volume offered	1,231,850 MMBtu
Total buy volume requested	508,000 MMBtu

Requested demand exceeds capacity by a factor of 5.5 and offered supply by 13.4. This market is rationed by the pipe rather than by the crossing of bids and asks.

Table 13. Clearing results for the synthetic Dahej book.

Result	**Value**
Ex-hub clearing price π	1,082.31 per MMBtu
Capacity shadow price μ	296.23 per MMBtu
Seller threshold τ	786.08 per MMBtu
Total welfare	33.04 million per day
Consumer surplus	5.19 million
Producer surplus	27.86 million
Of which scarcity rent	27.21 million
Buyers with a non-zero allocation	8 of 40

Result	Value
Sellers with a non-zero allocation	2 of 15
Volume cleared	91,850 MMBtu, capacity fully used

Exactly one order on each side is fractional, at 1,082.31 on the buy side and 786.08 on the sell side. These are precisely the clearing price and the seller threshold, which is the signature of marginality the case table predicts.

These results come from a synthetic order book. They are used to characterise the mechanism's behaviour under a plausible ratio of capacity to demand rather than to measure the real market.

9.5 The problem this exposed

The original approach was iterative: walk sorted buy and sell curves, match greedily, then apply the capacity limit. Moving that limit inside the optimisation was correct, and the results above are what a correct treatment produces.

They are also, as a market outcome, bad.

The binding constraint pulled the equilibrium back and pushed the price up. With throughput capped well below the crossing point, the market clears at the top of the demand curve and the bottom of the supply curve, and the gap between them opens as scarcity rent. Some 84 per cent of welfare accrues to producers, and almost all of that is scarcity rent rather than production efficiency. This is a measurement of how binding the pipeline is rather than a judgement about sellers.

Two sellers exhausted the point. The two cheapest offers filled the capacity between them, so no third seller could be reached at any price. On the buy side, 8 of 40 buyers cleared. A correct clearing engine, applied to a genuinely constrained single point, produced an outcome that is structurally concentrated, and that is exactly the outcome that keeps participants away.

This is the point at which the project changed direction, and it is the concrete form of the concession made in Section 2.3. Fixing the algorithm had not fixed the market, because the constraint was never algorithmic.

9.6 Approaches considered and rejected

Monte Carlo simulation over participant entry. The first response was to model how the price and cleared volume respond as the number and composition of participants varies, and use that to characterise which participant mixes produce acceptable outcomes. We rejected this as a design tool. It answers the question of which participants should be present for the market to behave well, which in practice becomes a filter, and the objective is to increase participation rather than to select for it. A mechanism that behaves well only on a curated set of participants has not solved the problem. Monte Carlo is retained as a validation instrument, since it can measure how the clearing price and the capacity dual respond to changes in the participant mix, which is useful for infrastructure appraisal.

Restricting admission to the exchange. Considered and rejected for the same reason.

Reconsidering what counts as a good allocation. This was more productive. The question is whether it is better to fill a few participants fully or many participants partially. Under a single delivery point with binding capacity these are close to equivalent in welfare terms, since the constraint binds either way, but they are not equivalent in participation terms. A buyer who is partially filled has

traded on the exchange. A buyer who was rejected has not, and next session may not return. This is what makes divisibility a market-design commitment rather than a modelling convenience.

Sourcing one order from several delivery points. The decisive idea. If a buyer's requirement can be met partly from one point and partly from another, then a single point's capacity ceases to be the binding limit on that buyer's participation. Capacity that is idle at one point can serve demand that is rationed at another. This does not relax the physical constraint. It uses more of the constraint that already exists. That is Section 10.

9.7 Ancillary results

From the clearing price to a delivered price. Because bids are ex-hub, settlement is a separate pass that never feeds back into the allocation. Each matched buyer pays the clearing price plus their transmission tariff, marketing margin where applicable, exchange fee and VAT. Sellers receive the clearing price less the exchange fee. Two consequences belong in the rulebook. The clearing price is identical for every participant at the point, so all variation in delivered price is attributable to identifiable third-party charges. And a buyer's delivered price may exceed their own limit price, which is not an error, since the limit was a bid for gas at the point. This should be surfaced at order entry.

Multi-seller sourcing. The formulation matches aggregate volume rather than counterparties, so a buyer's accepted volume may be drawn from several sellers. This is compatible with the anonymity of the exchange and arguably is what makes it meaningful. The buyer contracts with the exchange and the exchange with the seller, and no bilateral pairing is created or disclosed. Gas delivered ex-hub is fungible, so this raises no quality question at a single point. It becomes substantive in Section 10, where volume may be sourced from a different point.

10. Multiple Delivery Points

10.1 Approach

Section 9 clears one delivery point. This section asks what changes with fourteen, and the answer is less than expected. The objective is unchanged. The decision variables are unchanged except that each order carries a delivery-point label. One balance constraint becomes one per point per day, each with its own dual, so each point and day forms its own price from its own supply and demand.

The substantive question is whether points should be treated as interchangeable. They should not. Finding 3 of Section 3.1 shows materially different prices at different points on the same date for the same product, with gaps of 300 to 1,480 rupees persisting day after day. If gas moved freely between these points, arbitrage would close those gaps within a session. It does not. Each delivery point is therefore its own market, and a hub is a regional grouping used for reporting and indexing with no role in clearing. The balance constraint is imposed at the delivery point, never at the hub; points within one hub clear at materially different prices, which they could not do if balance were imposed at hub level.

10.2 What we deliberately do not model

No pipeline flows between points, no inter-point transport, and no capacity coupling.

The reason is institutional rather than mathematical. The exchange does not dispatch gas and holds no transmission capacity rights. Participants book transport separately with the transporter, ahead of the auction. Transport is a constraint on the participant, already resolved before they submit, since they bid only at points they can actually take gas from and their transport reality is already inside the price they bid.

This is where a gas exchange differs from a power exchange. In electricity the network is dispatched simultaneously with energy by the same operator, so flows are inseparable from clearing. On IGX transmission is booked separately and ahead. Importing a network model would mean the exchange scheduling flows nobody asked it to schedule, against capacity it does not hold.

A practical benefit follows, and it is larger than it looks. Because points do not interact, no constraint links the price at one point to the price at another, so the set of valid price systems separates into independent ranges, one per point and day. Prices may therefore be selected point by point, independently, and the rule of Section 7.5 applied to each. With flow variables the ranges would be tied together and selection would have to be made jointly across the network, a materially harder problem. Not modelling transport buys this simplification.

10.3 Method: multi-point orders

Participant flexibility across points is expressed through the order rather than through the network.

A participant able to take or make delivery at several points submits one order carrying a set of eligible points, a limit price for each, and a single aggregate quantity cap. The decision variable gains a point index, and one constraint caps the sum across all eligible points at the requested quantity. The acceptance condition is the single-point condition applied point by point with that point's own price, so

a buyer clears at a point only if their price for that point clears that point's market. The problem remains a linear programme.

The participant sets the price difference between points. That difference is their own transport economics, set by the party who knows their tariffs and booked capacity. The exchange holds no tariff table, assumes no distance relationship, and takes no position on what is economic for anyone.

This addresses the Section 9 finding directly. Today a participant wanting flexibility must submit separate orders at each point and risk being filled at both, ending up with twice the gas they wanted or gas they do not have. Most avoid the risk by picking one point and forgoing the others, which removes competition between points that should exist. Under a single order with an aggregate cap, the double-fill risk disappears, points compete for the same order, and thin points gain depth. Given that most points in the traded data are thin, this is where the participation gain lies.

Mhaskal has two local orders, so one buyer faces one seller and the seller clears only two-thirds of what it offered because there is no second buyer to reach. Making three additional orders eligible at the point raises the order count to five, clears 3,000 MMBtu instead of 2,000, and fills the local seller completely. Total welfare across both points rises by 110,000 rupees per day.

Two observations from that result are worth stating. The price at the thin point rose from 760 to 870, because the old price was low as a consequence of a single buyer facing a partially filled seller rather than of genuine competition. Depth moves a price toward what the gas is worth there, which need not be downward. And eligibility is optionality rather than relocation, since two of the newly eligible orders still cleared entirely at their original point because that is where they were better off. Declaring a point costs a participant nothing and commits them to nothing.

10.4 Challenges

Symmetry is not optional. If buyers could nominate several points while sellers could not, buyers would gain optionality that sellers lack. That is an asymmetry in the rules of participation rather than merely in prices, and it could not be defended. The construction must be written symmetrically for both sides.

Credit and position limits apply to the aggregate, not per point. Otherwise a participant could submit what is really one demand as several point-orders, exceeding their own limits or influencing several point prices at once.

Administered-price gas needs no special treatment. IGX already separates it by naming it as a distinct delivery point, as with Hazira-ONGC, Mhaskal (Non-HPHT) and Gadimoga (Ceiling Price). It is simply another point with its own balance and its own price, which means domestic gas can be included from the first version rather than deferred.

10.5 Scaling

The multi-point problem remains a linear programme, and because points do not interact it is the same calculation run once per point

11. What Kind of Market Is IGX Today

The preceding sections treat IGX as a market to be cleared well. This section asks a prior question: is it yet a market in the economic sense? The evidence suggests that it works more as a trading and settlement platform. A market discovers price through competition among many orders. The data suggests that at many points and on many days, IGX records prices formed in private deals.

11.1 The intuition of a market

A working market has four features. Many buyers and sellers meet in one place. Each order competes against all the others. The price comes out of that competition rather than out of a negotiation. And the price is trusted enough to be used as a reference elsewhere.

11.2 How IGX compares

Table 14. How IGX compares with the features of a functioning market.

Feature of a market	**What IGX shows today**
Many participants meet	Roughly 190 buyers face about 35 sellers, and many points are thin on a given day
Orders compete	76 per cent of lines have equal buy, sell and traded quantity, consistent with one-to-one matches
Price from competition	Runs of Daily contracts at one price across many days suggest pre-agreed packages
Price is a trusted reference	GIXI is published, but rests on thin and partly pre-arranged volumes

11.3 Defining the current market type

In industrial organisation terms, IGX today is best described as a thin, fragmented set of local markets, one per delivery point. Many points resemble a bilateral monopoly or a small supply-side oligopoly, with few sellers facing few buyers. Much of the trade is closer to over-the-counter dealing that uses the exchange to register and settle.

Three forces keep it there. Products do not match multi-day needs, so packages are agreed privately (Section 3). Delivery capacity at key points is tight, so a few sellers fill it (Section 9.5). And the tax treatment of gas raises the delivered price, which keeps smaller buyers away (Section 8).

11.4 Where it needs to be

The target is a liquid, competitive two-sided auction at each major delivery point. In that market many orders compete, a marginal order rather than a negotiation sets the price, and block packages clear inside the auction. Perfect competition is not a realistic aim. Workable competition is: enough depth that no single participant sets the price, and enough transparency that the price can be checked.

The engine set out in this paper is necessary for that market but not sufficient. A correct optimiser only optimises the orders it receives. The policy recommendations in Section 13 are what bring more orders in.

12. Limitations and Further Work

12.1 Delivery-point capacity

Whether points carry enforceable daily limits is unresolved. If they do not, the capacity constraint is dropped and the seller threshold τ collapses to the clearing price π, which changes a derived result rather than a parameter. Since the Section 9 findings turn on capacity binding hard, this needs confirmation from IGX before the engine is relied on.

12.2 Transport pricing

The engine prices ex-hub and holds no tariff table. Participants price the difference between delivery points themselves.

Two reporting layers are worth building regardless and require no model change: displaying delivered cost after clearing, and an indicative pre-trade point recommendation computed from the previous session's prices and the buyer's own tariffs. Any such recommendation is indicative only, since it uses prices from a session that has closed and the current session's prices depend on the orders in it, including the buyer's own.

Moving transport inside the objective is a different market design rather than an extension of this one. It would require the exchange to hold per-buyer, per-point tariff data it does not currently hold, and to take a position on transport it does not currently take. It also embeds a regulatory distortion, because transmission tariffs are zonally averaged and not necessarily cost-reflective for any particular flow, so any distortion in the tariff book would become a distortion in the clearing outcome. Whether the exchange should internalise transport at all is a question for IGX and PNGRB rather than an engineering choice.

12.3 Linepack and multi-period clearing

The balance constraint asserts there is no warehouse. Physically this is not quite true, since pipelines hold linepack, which is real short-term storage. Admitting it requires an inventory state variable and inter-period coupling, changing the problem from clearing a delivery period to clearing a horizon. This is substantial and should not be attempted until the current model is in production.

Linepack flexibility is also the physical basis for arguments in favour of the side-payment model this paper declines. If that commitment is ever revisited, this is the ground on which the argument will be made, and the ΔW series of Section 5.6 is the evidence that would inform it.

12.4 Sequential sessions and products that are not day-specific

Sessions are independent because a concluded IGX trade is a matched pair contributing equal and opposite quantities to the same point-day balance. If IGX introduced a product that is not day-specific, such as storage, swing, or a nomination right rather than a fixed daily quantity, delivery days would become linked and an earlier session could leave a genuine open position at a point-day. The engine would then need a committed-position term in the balance. It is recorded here so it does not have to be rediscovered later.

12.5 Lots and rounding

Orders are submitted in lots and half-lots do not trade, but the formulation treats volume as continuous, so a fractional fill can produce a non-lot quantity. Three options exist: round at settlement and absorb the residual, impose lot granularity as an integrality constraint, or set the lot size small enough that rounding is immaterial. The second converts every order to a binary decision for the same reasons as block orders and is not recommended. The first is recommended, with the rounding rule and the treatment of the residual written into the rulebook.

12.6 Gas quality

The formulation treats gas at a delivery point as fungible, so an accepted volume may be sourced from several sellers. Where calorific value or specification differs materially between sources, that assumption needs stating explicitly, or quality must enter as a constraint.

12.7 Other open items

Time priority among equal bids. Where orders share a limit price and only some can be filled, the formulation is indifferent. A tie-break must sit outside the optimisation and be stated in the rulebook. Time priority is the natural choice and is what the continuous book already uses.

Collateral adequacy. Orders should be validated against posted margin before entering the book rather than after clearing.

Contract types. The formulation clears one contract type at a time. Clearing several simultaneously requires separate solves or an explicit statement of their substitutability.

Solve-time budget. For a mixed-integer problem, an unfinished search must never be published as though it were optimal. If the budget is exceeded, the engine falls back to the best fully price-supported result found so far, recording the optimality gap, and if none exists the session is declared uncleared.

13. Conclusion and Policy Recommendations

An exchange must decide who trades, how much, and at what price. This paper has argued that answering the first two by constrained optimisation makes the third an implication rather than a decision, and that this is worth doing for reasons that are institutional as much as technical.

The technical case runs through the mathematics of Section 4. A clearing price obtained as the Lagrange multiplier on the balance constraint is the only price consistent with the allocation. It is created by the Lagrangian as a purely mathematical object, and the stationarity condition then reveals it as the quantity each participant's limit price must be compared against. The same multiplier read as a rate of change is the value of one more unit of gas at that point on that day, which is why it carries two names. The accept and reject rules a rulebook would otherwise have to assert are properties of that solution rather than additions to it. And quantities the exchange does not currently produce fall out of a solve it would be running anyway, chief among them the value of scarce delivery-point capacity.

All of that rests on convexity. A linear objective over a polyhedral feasible set gives strong duality, and strong duality is what guarantees that a supporting uniform price exists at all. This is easy to take for granted precisely because it never fails in the divisible case.

The institutional case is that the welfare-maximising allocation clears here, and here is the derivation, is a defensible sentence to a participant or a regulator, and the curves crossed there is a weaker one.

Extending to multiple delivery points is mathematically routine. The gains do not come from coupling the network, since the exchange neither dispatches gas nor holds transmission rights. They come from letting a single order be eligible at several points under one aggregate cap, which removes the double-fill risk currently forcing participants to pick one point and forgo the others, and brings competition to points that are thin today. The obstacles are not computational but sit in the rulebook, and they should be settled before the code is written.

The block order is where the mathematics stops being comfortable. Convexity is lost, strong duality fails, and no uniform price need support the outcome. The response set out in Section 5 is to refuse to read prices from a problem containing switches, to lock the switches and price the convex problem that remains, to test that no accepted package is out of the money across its profile, to reject branches rather than participants when that test fails, and to measure and publish the welfare given up to price uniformity. None of that is free, and the paper's position is that the cost should be measured rather than hidden.

13.1 Policy recommendations

Section 11 defined IGX today as a thin, fragmented set of point markets with much over-the-counter style dealing. The recommendations below are what would move it towards a competitive market. Each names who acts and how progress would be measured.

Table 15. Policy recommendations, who acts, and how progress is measured.

#	Recommendation	Who acts	How progress is measured
1	Introduce block (strip) orders covering any set of delivery days, accepted all or nothing	IGX, with PNGRB approval of the rulebook	Share of multi-day volume cleared in the auction rather than registered as runs of Daily contracts
2	Clear on welfare first and volume second, with a published price rule, tie-break and reproducibility standard	IGX	Every session re-derivable from stored inputs
3	Allow one order to be eligible at several delivery points under one quantity cap, equally for buyers and sellers	IGX	Number of orders per point at thin points
4	Publish daily capacity at each delivery point and its shadow value μ	IGX with transporters; PNGRB	μ used as an input in capacity expansion decisions
5	Publish the cost of uniform pricing (ΔW) and the list of paradoxically rejected blocks every session	IGX	A persistently large ΔW triggers a review of the no side payments rule
6	Flag pre-arranged trades separately from auction-cleared trades, and compute an auction-only version of GIXI	IGX; PNGRB	The gap between the two indices narrows as competition grows
7	Bring natural gas within GST so that tax on gas and tax on transport can be credited against each other	GST Council; states; Ministry of Petroleum and Natural Gas	Delivered cost gap between gas and competing fuels for small industrial buyers
8	Until then, publish a delivered cost breakdown alongside each clearing price	IGX	Breakdown available for every delivery point
9	Strengthen access to unused booked pipeline capacity, for example through capacity release or secondary trading	PNGRB; transporters	Share of booked capacity left idle at congested points

Recommendations 1 to 6 need no change in law and can be piloted at one or two major delivery points first, such as Dahej and Hazira. Recommendations 7 to 9 sit with policymakers, and the exchange can support them with data it already produces.

13.2 Toward a market worth entering

The narrower argument above serves a wider one. A market of roughly 190 buyers facing roughly 35 sellers, in which the largest packages are negotiated privately and registered afterwards, is not a market whose price means very much.

Each of the changes proposed here works on that directly. Divisible acceptance means a participant who cannot win an entire allocation is filled in part rather than rejected outright, and a buyer who has traded returns where a buyer who was rejected may not. Multi-point eligibility means thin points draw orders from beyond their immediate catchment. The block order means the largest transactions in the data can be submitted to the auction rather than agreed on the telephone. Prices that can be re-derived months later mean a participant leaving a negotiated price for an auction price can defend that decision internally. Each reduces a specific reason not to participate, and participation is the only thing that

eventually dissolves the asymmetry, because no clearing rule can manufacture competition that is not in the order book.

But the mechanism can only go so far, and it would be dishonest to claim otherwise. Two constraints sit outside it entirely.

The fiscal treatment. Because natural gas sits outside GST while its transportation does not, the cascade described in Section 8 falls on the buyer. For a buyer taking modest volumes the burden is disproportionate to the quantity, and that is exactly the buyer the exchange most needs to attract and exactly the one for whom gas loses against coal or liquid fuels at the burner. Bringing gas within GST, so that commodity and transportation are taxed under one creditable regime, would remove a deterrent no clearing algorithm can reach. Until then, the least the exchange can do is publish the decomposition, showing what the market formed and what was added afterwards, so the burden is visible to those in a position to act on it.

Delivery-point capacity. The Section 9 results show a market rationed by the pipe rather than by the crossing of bids and asks, with capacity binding hard enough that two sellers exhausted a point and most buyers cleared nothing. Under that constraint a correct clearing engine still produces a concentrated outcome, because the constraint rather than the algorithm is what binds. Capacity expansion is therefore not an adjacent infrastructure question but a precondition for the market this paper describes. The engine contributes something useful here, since the capacity shadow price is a defensible per-day valuation of additional throughput, produced as a by-product of ordinary clearing, and it is the correct object to place in front of an expansion decision.

13.3 What remains

Nothing in this paper creates liquidity by itself. It makes a missing product expressible, makes prices derivable rather than chosen, and makes the cost stack legible. Whether participants respond is a matter for them, for IGX, and for the policy changes that sit beyond the exchange's control.

14. References

Benders, J. F. (1962). Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4, 238–252. https://doi.org/10.1007/BF01386316

Bertsimas, D. and Tsitsiklis, J. N. (1997). *Introduction to Linear Optimization*. Athena Scientific.

Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press. https://doi.org/10.1017/CBO9780511804441

Codato, G. and Fischetti, M. (2006). Combinatorial Benders' cuts for mixed-integer linear programming. *Operations Research*, 54(4), 756–766. https://doi.org/10.1287/opre.1060.0286

da Silva, R. F. F. and Schouery, R. C. S. (2024). A branch-and-cut-and-price algorithm for cutting stock and related problems. arXiv:2308.03595v2 [math.OC]. https://doi.org/10.48550/arXiv.2308.03595

Garcia-Bertrand, R., Conejo, A. J. and Gabriel, S. A. (2006). Electricity market near-equilibrium under locational marginal pricing and minimum profit conditions. *European Journal of Operational Research*, 174(1), 457–479. https://doi.org/10.1016/j.ejor.2005.03.037

Gribik, P. R., Hogan, W. W. and Pope, S. L. (2007). Market-clearing electricity prices and energy uplift. Working paper, Harvard University.

Heather, P. Oxford Institute for Energy Studies papers on European traded gas hubs (editions to be specified).

Hogan, W. W. and Ring, B. J. (2003). On minimum-uplift pricing for electricity markets. Working paper, Harvard University.

Land, A. H. and Doig, A. G. (1960). An automatic method of solving discrete programming problems. *Econometrica*, 28(3), 497–520. https://doi.org/10.2307/1910129

Liberopoulos, G. and Andrianesis, P. (2016). Critical review of pricing schemes in markets with non-convex costs. *Operations Research*, 64(1), 17–31. https://doi.org/10.1287/opre.2015.1451

Madani, M. and Van Vyve, M. (2015a). Computationally efficient MIP formulation and algorithms for European day-ahead electricity market auctions. *European Journal of Operational Research*, 242(2), 580–593. https://doi.org/10.1016/j.ejor.2014.09.060

Madani, M. and Van Vyve, M. (2015b). A MIP framework for non-convex uniform price day-ahead electricity auctions. arXiv:1410.4468v2 [math.OC]. https://doi.org/10.48550/arXiv.1410.4468

Martin, A., Müller, J. C. and Pokutta, S. (2014). Strict linear prices in non-convex European day-ahead electricity markets. *Optimization Methods and Software*, 29(1), 189–221. https://doi.org/10.1080/10556788.2013.823544

Meeus, L., Verhaegen, K. and Belmans, R. (2009). Block order restrictions in combinatorial electric energy auctions. *European Journal of Operational Research*, 196(3), 1202–1206. https://doi.org/10.1016/j.ejor.2008.04.031

Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press. https://doi.org/10.1017/CBO9780511813825

NEMO Committee (2020). *EUPHEMIA Public Description: Single Price Coupling Algorithm*.

NEMO Committee (2025). *EUPHEMIA Public Description: Single Price Coupling Algorithm.* Version of 18 December 2025.

O'Neill, R. P., Sotkiewicz, P. M. and Rothkopf, M. H. (2007). Equilibrium prices in power exchanges with non-convex bids. Working paper, revised July 2007.

O'Neill, R. P., Sotkiewicz, P. M., Hobbs, B. F., Rothkopf, M. H. and Stewart, W. R. (2005). Efficient market-clearing prices in markets with nonconvexities. *European Journal of Operational Research*, 164(1), 269–285. https://doi.org/10.1016/j.ejor.2003.12.011

Padberg, M. and Rinaldi, G. (1991). A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems. *SIAM Review*, 33(1), 60–100. https://doi.org/10.1137/1033004

Petroleum and Natural Gas Regulatory Board (2020). *PNGRB (Gas Exchange) Regulations, 2020*.

Ruiz, C., Conejo, A. J. and Gabriel, S. A. (2012). Pricing non-convexities in an electricity pool. *IEEE Transactions on Power Systems*, 27(3), 1334–1342. https://doi.org/10.1109/TPWRS.2012.2184562

Samuelson, P. A. (1952). Spatial price equilibrium and linear programming. *American Economic Review*, 42(3), 283–303.

Savelli, I., Cornélusse, B., Giannitrapani, A., Paoletti, S. and Vicino, A. (2018). A new approach to electricity market clearing with uniform purchase price and curtailable block orders. *Applied Energy,* 226, 618–630. https://doi.org/10.1016/j.apenergy.2018.06.003

Scarf, H. E. (1994). The allocation of resources in the presence of indivisibilities. *Journal of Economic Perspectives*, 8(4), 111–128. https://doi.org/10.1257/jep.8.4.111



ASHOKA
UNIVERSITY

Campus Location
Plot No. 2, Rajiv Gandhi Education City, National Capital Region P.O. Rai, Sonepat Haryana-131029

Delhi Office
Ashoka University Office, 222, Second floor, Okhla Industrial Estate, Phase III, New Delhi-110020