

Price discrimination and limits to arbitrage in global LNG markets

EPRG Working Paper 1317

Cambridge Working Paper in Economics 1340

Robert A. Ritz

Abstract Gas prices around the world vary widely despite being connected by international trade of LNG. Some industry observers argue that major exporters (e.g., Qatar) have acted irrationally by failing to engage in price arbitrage. This is also difficult to reconcile with a perfectly competitive model in which price differences exist solely because of transport costs. We show that a model with market power can rationalize observed price differentials and trade flows. We highlight how different features of the LNG market limit the ability and/or incentive of other players to arbitrage, and discuss the potential impact of US LNG exports.

Keywords International trade, limits to arbitrage, LNG pricing, market power, natural gas, price discrimination

JEL Classification D40, F12, L95

Contact rar36@cam.ac.uk
Publication September 2013

Price discrimination and limits to arbitrage in global LNG markets

Robert A. Ritz*

Faculty of Economics & EPRG

University of Cambridge

`rar36@cam.ac.uk`

September 2013

Abstract

Gas prices around the world vary widely despite being connected by international trade of LNG. Some industry observers argue that major exporters (e.g., Qatar) have acted irrationally by failing to engage in price arbitrage. This is also difficult to reconcile with a perfectly competitive model in which price differences exist solely because of transport costs. We show that a model with market power can rationalize observed price differentials and trade flows. We highlight how different features of the LNG market limit the ability and/or incentive of other players to arbitrage, and discuss the potential impact of US LNG exports.

Keywords: International trade, limits to arbitrage, LNG pricing, market power, natural gas, price discrimination

JEL classifications: D40 (Market Structure and Pricing), F12 (Trade with Imperfect Competition), L95 (Gas Transportation)

*Thanks are due to Alberto Behar, Chi Kong Chyong, Simon Cowan, Philipp Koenig, Peter Neary, David Newbery, Pierre Noël, Robin Smale, Yves Smeers, John Ward and Katja Yafimava for helpful discussions and advice, to Yudi Wang and Yichi Zhang for research assistance, to Platts for providing some of the price data, and to participants at an Energy & Environment seminar in Cambridge for their feedback. All views expressed and any errors are my own.

1 Introduction

Not so long ago, there was a widespread conjecture that gas prices around the world would converge. The idea was that international trade in liquefied natural gas (LNG) would connect previously separate geographies—notably the regional markets of Asia, Europe and the US—and thereby link their pricing.

The volume of LNG trade has indeed grown significantly since the early 2000s, against the backdrop of liberalization of electricity and gas industries, and a general increase in demand for natural gas. Global investment in LNG infrastructure—liquefaction and regasification capacity—has increased, and such facilities are now spread across more countries. The global LNG tanker fleet has expanded significantly, and transport costs have fallen.¹

Importantly, too, contracting arrangements between buyers and sellers have become more flexible. Traditionally, an LNG project involves a bilateral long-term contract, say of a 20-year duration, between a buyer and seller, to back up the initial investment. However, there is an ongoing shift towards trade in spot and short-term markets; these have increased ten-fold since 2000, and now make up 25% of total LNG sales (GIIGNL, 2012).² This development has been aided, amongst other things, by the adoption of Master Sales Agreements for LNG, which create standardization and reduce transaction costs.

Yet gas prices around the world today vary widely, and these differences have become more pronounced in the aftermath of the Fukushima accident of March 2011. The average price of natural gas in 2012 was roughly US\$16/MMBtu in Japan, around \$9 in European markets like the UK, but only \$3 in the US. Some expect large regional price disparities to persist, including the International Energy Agency in its modelling scenarios for both 2020 and 2035. In short, the gas market today appears far from global.

For the case of the US, the reasons for price divergence are quite clear. First, the large-scale emergence of shale gas over the last few years has put strong downward pressure on US natural gas prices. Second, the US at present only has very limited LNG export capability; its infrastructure still reflects the assumption of the 2000s that the US would become a major LNG importer. As a result, the US market has been largely isolated from the rest of the world.³

¹Useful overviews of the LNG industry as of the mid 2000s are provided by Jensen (2003) and Yergin and Stoppard (2003), with a focus, respectively, on economics and geopolitics.

²Brito and Hartley (2007) argue that the expectation of a shift towards short-term multilateral trading can have self-fulfilling properties.

³Several applications to create LNG liquefaction facilities are currently pending US regulatory/political approval.

The other price gaps require a different explanation, and some industry observers have argued they imply that LNG exporters have been behaving “irrationally”. Several major LNG exporters make short-term sales to Asia but simultaneously supply Northwest Europe at far lower prices. This behaviour may appear irrational in that it entails a forgone profit = $|\text{price differential}| \times \text{quantity sold to the lower-priced market}$. For Qatari short-term sales to the UK, rather than to Japan, some estimates suggest a forgone profit of up to \$100 million per day (in late 2011), and a cumulative figure in the billions over the period from April 2011 to April 2012.⁴ LNG producers are, apparently, failing to engage in price arbitrage by not exiting the European market (at least for short-term sales).

The most immediate explanation for price divergence lies in transport costs. In particular, a simple perfectly competitive model predicts that the price difference across two regions served by an exporter equals the difference in the associated transport costs. Put differently, the “netback”—that is, price minus transport cost—for the exporter should be the same for each region.

The problem with this theory is that it cannot explain the kinds of price differences recently observed in LNG markets. Consider again the case of Qatar, which is the largest LNG producer (with a global market share of around 30%). Figure 1 shows the differential between the Platts JKM (Japan Korea Marker) price for Asian LNG and the UK NBP (National Balancing Point) price, plotted against the difference in transport costs between shipping from Qatar to Japan and Qatar to the UK, respectively.^{5,6} Prices are up to \$10/MMBtu higher in Asia than in the UK, while the corresponding transport costs are approximately identical. Perfect competition, by contrast, predicts that these two differentials should coincide.⁷

The simple theory also fails to explain the data on two broader counts. First, over large parts of the sample period (early 2010 until early 2013) it even predicts the wrong sign: transport costs to the UK are typically slightly *higher* than to Japan, while prices are much lower. Second, there is “excess volatility”: transport costs are far too stable to be able to explain the observed volatility of (relative) gas prices.

The perfectly competitive model cannot account for other producers’ behaviour. For example, Peru is in a similar position to Qatar in that its transport-cost differential to European and Asian markets is usually very small (Platts, 2012), and yet

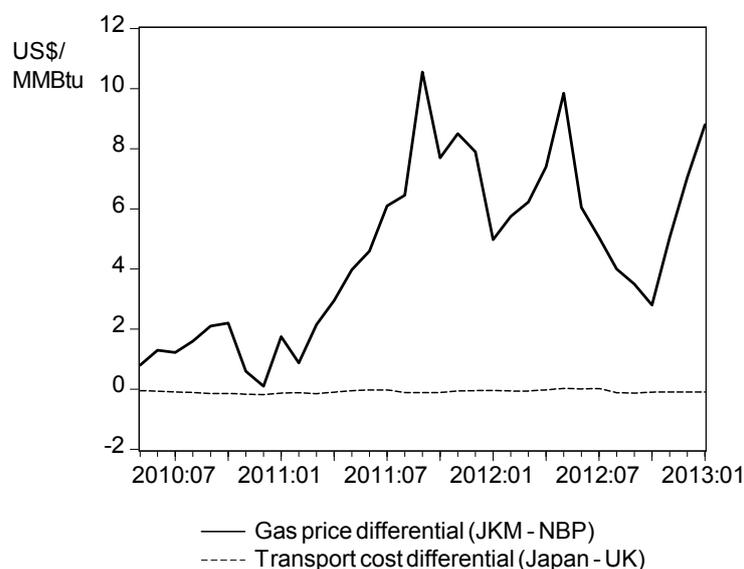
⁴These deliveries are estimated as 75% higher than contractual obligations.

⁵Qatari sales to Japan and UK are the two largest routes of short-term LNG, with global shares of 10% and 6% respectively (GIIGNL, 2012).

⁶Source: Calculations based on data from Platts, Poten & Partners, and ICAP (via Bloomberg).

⁷This theoretical prediction should, of course, not be taken too literally. Temporary deviations are to be expected in the face of short-term demand and supply shocks. However, large price divergences persisting for several years are difficult to reconcile with this theory.

Figure 1: Qatar LNG sales to Japan vs UK—Differences in gas prices compared to differences in transport costs



it still makes short-term sales to both (GIIGNL, 2012). For other major producers, such as Nigeria and Trinidad & Tobago, transport costs to Asia are indeed higher (by around \$2–3.50/MMBtu) but still often not sufficient to explain observed sales and prices.

In this paper, we instead suggest that regional price differentials can arise because of LNG exporters’ market power. Consider a producer who can sell uncommitted LNG into two export markets. In general, profit-maximization implies that the producer equalizes marginal revenue, net of the marginal cost of production and transport costs, across the two markets. If transport costs are identical—as is roughly the case for Qatar’s sales to the UK and Japan—then export quantities are such that marginal revenues for each region are equal.

Our key point is simply that equalizing marginal revenue is not necessarily the same thing as equalizing price. Put differently, for an exporter with market power, the “arbitrage” process stops when its marginal revenues are equalized; it is entirely possible that this optimally leaves prices across markets far apart. This basic argument extends straightforwardly to a producer selling into more than two export markets, and to situations with a capacity constraint on production. Moreover, the argument does not depend on any particular assumption on the mode of (strategic) competition in LNG; it applies for a monopolist, a Cournot-Nash player, a dominant firm facing a competitive fringe, and so on.⁸

⁸We do not suggest that LNG producers are colluding, but rather that at least some of them

We provide a new general formula for relative prices between any two export markets, in terms of a seller’s transport costs and price elasticities of demand.⁹ Incorporating market power can rationalize recent gas prices and trade flows by tracing them to relative demand conditions. The Fukushima accident, for instance, effectively switched off large parts of Japanese nuclear power, leading to an increase in demand for imported LNG to “fill the gap”. From the viewpoint of an individual LNG seller (with a degree of market power), under fairly general conditions, an upward shift in market demand translates into a lower demand elasticity—and thus a higher price. Importantly, note that local demand conditions play no role in the competitive model, in which price differences are *solely* driven by transport costs.

We can also offer a perspective on the possibility that the US will become a large LNG exporter over the coming years. What is the likely impact on gas prices? Our analysis makes clear that US and non-US prices will not necessarily converge as a result, even allowing for transport costs. It also suggests that any model of the effect of US LNG exports is likely to be incomplete if it does not take market power into account. For example, a recent model-based simulation for the US Department of Energy incorporates general-equilibrium effects but assumes that LNG producers do not respond strategically to US market entry.¹⁰

Our model of price discrimination with imperfect competition is static and thus does not capture intertemporal features of the LNG market (such as storage).¹¹ In practice, there may be a dynamic interaction between short-term LNG prices and long-term contracts. (Our model examines pricing incentives in short-term markets, taking long-term commitments as given.) In particular, a large proportion of LNG imports is still governed by long-term contracts whose terms may be renegotiated from time to time. Some LNG producers may have been reluctant to push down short-term prices in Asia insofar as this would make it more difficult to sustain “high” prices on long-term contracts in the future. At its core, this argument has a similar flavour to ours; while our model is based on exporters’ market power in short-term markets, this argument is essentially about exporters’ bargaining power in negotiations of long-term contracts.

Related literature. A number of empirical papers have examined price conver-

have a degree of market power (i.e., are not textbook price-takers) in some of their export markets. See Egging, Holz, von Hirschhausen and Gabriel (2009) for a recent analysis of “Gas OPEC”.

⁹Our approach does not rely on production cost data.

¹⁰NERA (2012) assumes Qatar is “large” but does not alter its production strategy in response to US exports, while other producers are represented as a competitive fringe; the model is augmented with (exogenous) mark-up adjustments in order to be able to replicate observed regional gas prices.

¹¹Chaton and Durant-Viel (2013) analyze the value of gas storage when firms have market power.

gence in natural gas markets.¹² Siliverstovsa et al. (2005) obtain mixed results from a cointegration analysis using data from the early 1990s until 2004; they find evidence for market integration between European and Japanese markets, but no integration between North America and Japan. Over the period 1999 to 2008, Neumann (2009) finds increased convergence of gas spot prices between North America and Asia. However, there does not appear to be any econometric analysis of gas-price convergence post-2011.¹³ Others have focused on price convergence within regional markets; for example, Doane and Spulber (1994) employ similar techniques to an integrated, national market for natural gas in the US.¹⁴ By contrast, we here offer an economic-theory perspective on price non-convergence, with a view to explaining observed prices since March 2011. We use the standard non-cooperative approach to pricing strategies; see Hubert and Ikkonikova (2011) for a recent application of cooperative game theory to bargaining and investment in natural gas markets.

Also related is the literature on market power in crude oil, which estimates the level of market competitiveness with a particular emphasis on the role of OPEC, see, e.g., Salant (1976), Almoguera, Douglas and Herrera (2011), and Nakov and Nuño (2013). Much of this literature concludes that the world crude oil market can be modelled either as a set of dominant producers facing a competitive fringe, or as something close to Cournot-Nash competition. Thereby, most oil models simply assume a single global oil price; indeed, international price differentials in crude oil are typically small and mainly reflect quality differences between oil varieties from different regions. In contrast, this paper focuses on natural gas, for which the existing literature—and the market—is much less well-developed; it obtains results that apply for a wide range of competitive conditions, and focuses explicitly on the limits to international price arbitrage in LNG.¹⁵

The remainder of this paper proceeds as follows. Section 2 presents a model of LNG pricing across different export markets to formalize the market-power argument we have outlined above. Section 3 provides a discussion of how different features of the LNG market limit the ability and/or incentive of other players (such as LNG buyers and third-party traders) to engage in price arbitrage, including constraints in LNG shipping. Section 4 offers concluding remarks, including a discussion of the

¹²The chapters in Stern (2012) provide a detailed overview of how gas-pricing mechanisms vary across different regions.

¹³We conjecture that international price correlations have declined significantly (compared to the mid/late 2000s), and that the Fukushima accident represents a structural break.

¹⁴There are also several recent simulation-based models of gas-market integration in the EU, see, e.g., Holz, von Hirschhausen and Kemfert (2008) and Lise and Hobbs (2009).

¹⁵There may be an interesting analogy between the globalization of oil markets during the 1970s/1980s and ongoing developments in natural gas, but this is beyond our current scope.

potential effects of greater price arbitrage, in the future, on LNG prices, industry profits, and social welfare.

2 A profit-maximizing LNG exporter

Consider an LNG producer selling output into $N \geq 2$ export markets.

Let $p_i^k(x_i^k, y_i^k, X_i^{-k}, Y_i^{-k}; \theta_i)$ denote the inverse demand function producer k faces in market i ($i = 1, 2, \dots, N$), where p_i^k is the “spot” price of LNG,¹⁶ x_i^k is the quantity sold by the producer in the short-term market, while y_i^k is the quantity the producer has pre-committed to sell in market i by way of long-term contracts (which we here take as given). Analogously, X_i^{-k} is the vector of outputs sold by other producers in the spot market while Y_i^{-k} captures their long-term commitments.¹⁷ Other factors that affect demand in market i are summarized by the vector θ_i . For example, this might include the prices of coal, oil and other substitute products, the state of business cycle, other demand shocks, the weather, and so on.¹⁸

The producer’s cost function $C^k(\sum_{i=1}^N (x_i^k + y_i^k))$ depends on the sum of total quantities sold in all N export markets, including spot market sales as well as long-term commitments. Production may be subject to a capacity constraint such that $\sum_{i=1}^N (x_i^k + y_i^k) \leq \bar{Q}^k$.¹⁹ In addition to this, the producer incurs a transport cost t_i^k per unit of output sold to market i . This mainly reflects the cost of shipping and may vary across export markets depending on distance and other factors (and also vary across different producers).

Producer k ’s profit-maximization problem is to choose the amount of LNG to export to each market, given any long-term commitments already entered into:

$$\max_{\{x_i^k\}_{i=1}^N} \Pi^k = \sum_{i=1}^N p_i^k x_i^k - C^k(\sum_{i=1}^N (x_i^k + y_i^k)) - \sum_{i=1}^N t_i^k x_i^k$$

$$\text{subject to } \sum_{i=1}^N (x_i^k + y_i^k) \leq \bar{Q}^k.$$

We assume, without much loss of additional economic insight, that this problem is well-behaved with an interior solution for each of the N export markets. The Lagrangean for constrained optimization can be written as $L^k = \Pi^k + \lambda^k(\bar{Q} -$

¹⁶We allow prices to be producer-specific, e.g., to be able to reflect (small) quality differentials.

¹⁷From a buyer’s point of view, LNG from short-term markets and long-term contracts may be imperfect substitutes.

¹⁸Some of these factors may affect individual producers in different ways, and some may influence demand conditions in several export markets.

¹⁹Adding a production constraint in terms of minimum throughput to ensure smooth operation, $\sum_{i=1}^N (x_i^k + y_i^k) \geq \underline{Q}^k$, would not affect the following results.

$\sum_{i=1}^N(x_i^k + y_i^k)$), where $\lambda^k \geq 0$ is the shadow value of the capacity constraint (i.e., the value of an incremental relaxation of the capacity constraint), which is non-zero if the producer is capacity-constrained and zero if it is not).

The optimal output choice \hat{x}_i^k by producer k in market i satisfies the first-order condition $MR_i^k - MC^k - t_i^k - \lambda^k = 0$, where MR_i^k is marginal revenue from short-term sales, and MC^k is the marginal cost of production. Using the first-order conditions for any two export markets, say i and j , shows that these are related by

$$MR_i^k - t_i^k = MR_j^k - t_j^k.$$

This is the fundamental condition for profit-maximization. The cost of an additional unit of output is the same regardless where it ends up being sold in, both in terms of the marginal cost of production and the shadow value of using the capacity elsewhere. (This holds regardless of whether or not the producer is, in fact, capacity-constrained.) To maximize profits, therefore, the producer balances at the margin the contribution of each export market in terms of sales revenue and transport costs. So marginal revenue net of transport costs is equalized across export markets.

Marginal revenue in market i can be written as $MR_i^k = p_i^k (1 - 1/\eta_i^k)$, where η_i^k is the own-price elasticity of producer k 's demand. From here on, this elasticity is understood to be evaluated at producer k 's optimally chosen output \hat{x}_i^k , as well as at the *actual* levels of short-term output of other producers and corresponding actual long-term commitments, that is, $\eta_i^k = \eta_i^k(\hat{x}_i^k, y_i^k, X_i^{-k}, Y_i^{-k}; \theta_i)$.²⁰ Combining this expression with the fundamental condition for profit-maximization leads to:

Proposition 1 *A profit-maximizing producer k sells into $N \geq 2$ export markets with a common marginal cost (and possibly subject to a capacity constraint, \bar{Q}^k).*

(A) *In any two markets i and j , profit-maximizing prices (p_i^k, p_j^k) satisfy*

$$\frac{(p_i^k - p_j^k)}{p_i^k} = \frac{\eta_i^k}{(\eta_i^k - 1)} \left[\left(\frac{1}{\eta_i^k} - \frac{1}{\eta_j^k} \right) + \frac{(t_i^k - t_j^k)}{p_i^k} \right],$$

where (t_i^k, t_j^k) are transport costs and (η_i^k, η_j^k) are own-price elasticities of demand;

(B) *Any observed prices (p_i^k, p_j^k) and transport costs (t_i^k, t_j^k) in markets i and j can be rationalized by some values for the price elasticities of demand (η_i^k, η_j^k) .*

Understanding profit-maximizing prices. The formula for relative prices from (A) is rather general: it does not rely on any specific functional-form assumptions

²⁰A necessary condition for profit-maximization is that producer demand remains price-elastic in each market, $\eta_i^k > 1$. (Otherwise the producer could profitably reduce output.) Market-level demand elasticities can be significantly lower.

on demand and cost functions (e.g., linear, constant-elasticity, etc.), or on a particular form of competitive conduct in each export markets. Commonly-used models, e.g., perfect competition, monopoly, Cournot-Nash oligopoly, dominant firm with a competitive fringe, etc., are nested as special cases; the mode of competition may also differ across export markets. An informational advantage is that it does not feature production costs. (The result does also not assume that either consumers or other producers are payoff-maximizers; their behaviour, rational or otherwise, is fully captured by the producer's own-price elasticities of demand across export markets.)

To understand the properties of the model, consider a few special cases:

First, the simple perfectly competitive model is nested where the producer's demand elasticity $\eta_i^k \rightarrow \infty$ in each export market $i = 1, 2, \dots, N$. This corresponds to a situation in which the producer is a price taker without any market power (so its marginal revenue is equal to the market price in each market). Prices in any two markets satisfy $(p_i^k - p_j^k) = (t_i^k - t_j^k)$, and netbacks are equalized.

Second, suppose that the equilibrium values of the price elasticities of demand in two exports market are identical, $\eta_i^k = \eta_j^k = \hat{\eta} < \infty$. Then the expression for the price differential becomes $(p_i^k - p_j^k) = (t_i^k - t_j^k) \hat{\eta} / (\hat{\eta} - 1)$. Relative to perfect competition, (symmetric) market power thus exacerbates any price differential across export markets that is due to transport costs.

Third, assume that transport costs to two markets are identical, $t_i^k = t_j^k$. Relative price then satisfies $(p_i^k - p_j^k) / p_i^k = (\eta_j^k - \eta_i^k) / (\eta_i^k - 1) \eta_j^k$ so $p_i^k > p_j^k$ if and only if $\eta_i^k < \eta_j^k$. This shows that (i) prices can diverge across markets for reasons of market power, not transport costs, and that (ii) "stronger" markets, in which a producer faces a lower price elasticity of demand, have higher prices.

Fourth, if price elasticities and transport costs across two markets satisfy $\eta_i^k \leq \eta_j^k$ and $t_i^k \geq t_j^k$, then prices must satisfy $p_i^k \geq p_j^k$. Intuitively, market i is "far-and-strong", with greater market power as well as higher costs, while market j is "near-and-weak". If either of these relationships is strict, then $p_i^k > p_j^k$.

A model with market power can thus explain a far greater range of observed prices than the simple competitive model. Most importantly, transport costs are no longer the sole driver of price differentials; relative demand conditions across export markets now also play a key role. It relaxes the strong restriction that $\text{sign}(p_i^k - p_j^k) = \text{sign}(t_i^k - t_j^k)$, and also features "excess volatility" in prices, by going beyond the implication from the competitive model that $\text{var}(p_i^k - p_j^k) = \text{var}(t_i^k - t_j^k)$.

Rationalizing observed prices. Part (B) of the result is that the model with market power can rationalize *any* observed price differences between export markets.

The reason is as follows. Appropriate choice of the price elasticity for market i can generate any non-negative price-cost margin, ranging from zero (when $\eta_i^k \rightarrow \infty$) to arbitrarily large (when $\eta_i^k \rightarrow 1$), regardless of the underlying details of the producer's costs.²¹ So it is always possible to find an elasticity to rationalize the observed price in any market, and thus also to generate correct relative prices across markets.

Consider a numerical example based on Qatari LNG sales to Japan and Northwest Europe ($i = \text{Japan}$, $j = \text{UK}$, $k = \text{Qatar}$). Let prices $p_i^k = 16$ and $p_j^k = 9$, and assume, for simplicity, that transport costs are identical, $t_i^k = t_j^k$. It is not difficult to check that these relative prices can be rationalized by a pair of elasticities $\eta_i^k = 2$ and $\eta_j^k = 9$. Such prices can thus be explained by producers having relatively greater market power in Japan.

So recent claims that LNG producers are acting irrationally by simultaneously selling short-term cargoes to both Northwest Europe and Asian markets are not necessarily correct. It can be entirely rational for a profit-maximizing seller to pursue a strategy that leaves prices in Japan far higher, in response to stronger demand. In effect, it uses sales to the UK to keep prices in Japan high.²²

Why might producers have greater market power in Japan? The Fukushima accident effectively switched off large parts of Japanese nuclear power, leading to an increase in demand for imported LNG so as to “fill the gap”. From the viewpoint of an individual LNG seller (with a degree of market power), under fairly general conditions, an upward shift in market demand (captured formally by a change in θ_i) translates into a lower price elasticity of demand. This, in turn, typically leads to an increase in quantity of LNG supplied but also to an increase in its price, as is consistent with market experience since Fukushima.

More generally, it is frequently suggested that Asian buyers are more concerned about “security of supply” than European buyers. This translates into a higher willingness-to-pay for a unit of LNG and, all else equal, a lower elasticity. Furthermore, Asian buyers have fewer possibilities to substitute for LNG, notably because of more limited access to Russian pipeline gas.

Estimating producer-specific elasticities of demand. A feature of the model is that the pair of elasticities to rationalize the data is, in general, not unique. In the numerical example for Qatar, the data $p_i^k = 16$, $p_j^k = 9$, and $t_i^k = t_j^k$ are rationalized for *any* pair of elasticities (η_i^k, η_j^k) that satisfies $(1 - \eta_i^k/\eta_j^k)/(\eta_i^k - 1) = \frac{7}{9}$. Setting $(\eta_i^k, \eta_j^k) = (2, 9)$ is but one solution. Loosely put, getting the relative elasticities

²¹A negative price-cost margin cannot be profit-maximizing in the present model.

²²It is also economically inefficient in that different consumers are paying different prices for essentially the same good (so their marginal utilities are unequal).

across markets correct matters more than their absolute values.

Pinning down unique values for the elasticities requires more information and/or additional modelling assumptions. We here suggest three possible approaches:

First, recall that the above results did not rely on production cost information, or on any knowledge of whether the producer is, in fact, capacity-constrained. However, if such information is available, this immediately identifies the producer-specific demand elasticity. To see this, rewrite the first-order condition for profit-maximization in market i , using the relationship $MR_i^k = p_i^k (1 - 1/\eta_i^k)$, to obtain $\eta_i^k = p_i^k / [(p_i^k - t_i^k) - (MC^k + \lambda^k)]$. In the numerical example for Qatar, knowledge that $(MC^k + \lambda^k) = 6$ would select the values $(\eta_i^k, \eta_j^k) = (2, 9)$ as the unique way of rationalizing the data. Perhaps most realistically, an assessment that the producer is not capacity-constrained, i.e., $\lambda^k = 0$, together with data on (p_i^k, t_i^k, MC^k) identifies η_i^k for each individual market.²³

Second, it is possible, in principle, to estimate the elasticity for an individual markets, η_i^k , or the entire vector of elasticities $\{\eta_i^k\}_{i=1}^N$, using econometric techniques. This would require time-series data on prices and quantities in each market of interest, as well as relevant control variables (probably including data on producers' long-term contractual commitments).

Third, it may be possible to justify more specific assumptions on competitive conduct. For instance, it is quite common in the analysis of natural gas markets to assume Cournot-Nash competition between sellers. Then an individual producer k 's demand elasticity $\eta_i^k = \eta_i / s_i^k$, where η_i is the price elasticity of *market* demand, and $s_i^k \in (0, 1)$ is the producer's market share (in the appropriately defined market i). Market share data are generally easier to obtain, and it is usually easier to estimate a market-level elasticity than a producer-specific elasticity.

To continue our numerical example from above, suppose it is estimated that two markets have identical price elasticities $\eta_i = \eta_j = \frac{1}{2}$ (say for natural gas); the producer-specific elasticities $(\eta_i^k, \eta_j^k) = (2, 9)$ would then be generated by producer k having market shares $(s_i^k, s_j^k) = (25\%, 5\frac{5}{9}\%)$ respectively in the two markets. Alternatively, if the producer had identical market shares $s_i^k = s_j^k = 10\%$, then the corresponding market-level elasticities would be $(\eta_i, \eta_j) = (\frac{1}{5}, \frac{9}{10})$.

Unfortunately, we currently do not have sufficiently rich data on costs, prices, and quantities to be able to pursue either of the first two of these approaches; we return to the impact of price discrimination under different competitive conditions in our concluding remarks below.

²³Put differently, for each seller, there are N first-order conditions but $N + 2$ unknowns, namely, $\{\eta_i^k\}_{i=1}^N$, MC^k , and λ^k , but since $MC^k + \lambda^k$ is a sufficient statistic for (MC^k, λ^k) the system boils down to $N + 1$ unknowns.

3 Limits to price arbitrage in LNG markets

Our model offers an explanation for why it may be optimal for prices across regions to be different from the viewpoint of an LNG producer. Implicit in the model, as in virtually all literature on price discrimination, is that other players do not undermine this sales strategy. We here discuss a number of reasons, many particular to LNG markets, that either limit the ability of other players to engage in arbitrage or create incentives that work against pursuing arbitrage in the first place.

The textbook assumption that sustains price discrimination by a seller is that buyers cannot engage in resale. Although there is a trend towards more flexible LNG contracting arrangements, some “destination restrictions” appear to persist. For example, it is said that state-controlled LNG exporters still normally restrict the resale of their exports in a way that prevents them from being traded on commodity exchanges. This means that some price arbitrage opportunities, if they exist, cannot be exploited for contractual reasons.

LNG arbitrage may also be difficult because of limited shipping capacity.²⁴ Although there now are on the order of 400 vessels for transporting LNG, only a small proportion of the fleet is uncommitted, in the sense of not being tied to a long-term sales contract. Thus only few companies appear to have direct access to both uncommitted gas supplies and uncommitted LNG tankers. So an LNG buyer wishing to engage in price arbitrage may find that the shipping market is either unable or unwilling to provide transport at the required price.²⁵ Note that this latter argument has some similarity to our model from above; it involves market power in the shipping market rather than (or in addition to) the LNG market itself.²⁶

In addition to this, there are at least two reasons to do with vertical structure for why arbitrage, even if possible, may not be in the interest of an LNG buyer. First, while redirecting cargo, say, from northwest Europe to Japan may promise a higher price, it also means that the LNG buyer can no longer sell or use the gas further downstream in the European market. So redirecting cargo may also forgo downstream surplus, which works against the incentive to arbitrage. Second, ownership arrangements along the LNG supply chain are much more complex than in any

²⁴Our model assumes that an LNG producer takes transport costs as given when choosing its sales strategy, and that transport is available for any desired export volume to any market. Moreover, our numerical illustrations use publicly available shipping rates as being representative.

²⁵In some cases, there may also be compatibility issues; not all import terminals are able to receive deliveries from all types of LNG tankers.

²⁶There is also a potential feedback effect: Since short-term LNG typically involves longer distances than trade from long-term contracts, more price arbitrage tends to further tighten the shipping market, and may thus to some extent undermine itself. Thanks to Philipp Koenig for this point.

simple model. Several LNG players hold partial (<100%) ownership stakes at various points along the supply chain—including in LNG production and liquefaction, shipping, regasification, and downstream gas—as well as across different countries. Put simply, a company may be an LNG seller in country A, an LNG buyer in country B, and have an infrastructure stake in country C. Such a player’s overall profit function—and hence incentive structure—is more difficult to work out. However, it seems plausible that the overall incentive sometimes works against arbitrage. In any case, vertical issues mean that a simple comparison of netbacks may not be enough.

If neither LNG sellers nor LNG buyers have a strong incentive to engage in price arbitrage, what about third parties such as traders? A recent industry report offers an interesting perspective on this question: “The entry barriers to LNG trading are surprisingly high—new entrants require more than just experienced traders and trading systems. They must have access to cargoes, but the market’s liquidity is typically held captive by the LNG liquefaction owners/upstream suppliers who are understandably very reluctant to release volumes for traders to trade with. Traders must also have access to shipping, either via owned vessels or the charter market. Furthermore, certain ships can unload at certain terminals (e.g., many import terminals cannot accommodate Q-Max vessels). This can make it even more difficult to efficiently connect volumes to buyers.” (JP Morgan Cazenove, 2012). That is, physical arbitrage requires sufficient capacities along the entire supply chain; almost by definition, this is more difficult for third parties to secure.

It is worth highlighting a few other considerations which, in practice, make LNG arbitrage difficult and financially risky—and are typically neglected in models of price discrimination. The first is units. While the flow of gas is, in some sense, continuous, the economics of LNG transport involves an indivisibility: the unit of account is, in effect, a tanker. As a result, only players with sufficiently “deep pockets” can enter the market. The second is time. It can take two weeks, for example, to ship LNG from Qatar to Japan. Given the volatility of gas prices, it is possible for there to be a significant shift in relative prices over such a period of time. So risk management becomes an important factor, both for LNG sales and potential arbitrage activity. Although financial instruments for natural gas exist, the derivatives market specifically for LNG is relatively underdeveloped at present. Financial arbitrage, in general, can also be affected by the existence of agency costs and capital constraints (Shleifer and Vishny, 1997).

Finally, the extent of price arbitrage in international gas markets may be limited because arbitrageurs themselves have a degree of market power. This can result from a combination of the “lumpiness” of LNG trade and barriers to entry discussed

above. In such cases, the optimal way to exploit a profitable trading opportunity does not lead to prices being equalized, precisely because the arbitrageur realizes that her actions have a non-zero effect on prices.²⁷ For example, the optimal arbitrage strategy $\hat{\alpha} \equiv \arg \max_{\alpha} \{\alpha [p_i(-\alpha) - p_j(\alpha)]\}$ of buying α units in low-price market j to sell in high-price market i leaves $p_i(-\hat{\alpha}) \neq p_j(\hat{\alpha})$ whenever the arbitrageur has market power in at least one of the two markets.

Taken together, these arguments suggests that, there are significant limits to the scale of arbitrage activity which mean that gas price differentials—perhaps due to market power in the LNG supply chain—have persisted.

4 Concluding remarks

Despite being connected by international trade in LNG, gas prices around the world vary widely. It is particularly surprising that large price differentials have persisted for several years now, notably since the Fukushima accident of March 2011. Some industry observers have thus claimed that LNG producers are behaving irrationally by failing to engage in international price arbitrage. Such relative prices are also difficult to reconcile with a perfectly competitive model in which price differences arise solely due to transport costs.

This paper presents the first attempt in the literature to address this puzzle. It shows that observed prices and trade flows can be explained by LNG producers having a degree of market power. Arbitrage by a profit-maximizing exporter takes place by comparing marginal revenue across markets rather than only price. Differences in local demand conditions can leave prices far apart. We have argued that, in addition, a combination of incentives, market power, and other constraints tends to work against international arbitrage by LNG buyers and third-party traders.

So is gas a global market? This is partly a matter of definition. Yes, in the sense that several LNG exporters sell into almost all major markets (except the US), and thus connect their pricing—albeit imperfectly. No, in that there is currently no clear tendency towards a single uniform gas price (even adjusted for transport costs).

Looking ahead, a number of recent developments, on balance, suggest that the gas market may become (even) more global. Significant low-cost capacity may emerge in form of LNG exports based on US shale gas. Yet other LNG projects, notably in Australia, have higher-than-projected costs which may dampen future

²⁷See also Borenstein, Bushnell, Knittel and Wolfram (2008) on pricing in Californian electricity markets around the time of the Enron collapse; they show that arbitrage opportunities existed between spot and forward markets but suggest these were left unexploited due to a combination of market power and arbitrageurs' fear of regulatory penalties.

supply. Production is also becoming more flexible; floating liquefaction plant and tankers with onboard regasification capabilities should make output more responsive to relative prices, and recent plans (e.g., in Japan) to introduce LNG futures contracts would facilitate hedging and arbitrage.

A natural question therefore is, how will greater price arbitrage affect the global LNG market? The existing theoretical literature on third-degree price discrimination offers some partial answers.²⁸ It focuses on the effect of moving from “uniform pricing”, where firms are forced (e.g., by regulation) to set identical prices in all markets, to price discrimination, where there are no such constraints on relative prices. Turned on its head, it therefore addresses the impact of an extreme scenario: moving from unconstrained price discrimination by LNG exporters to a world with perfect, costless arbitrage and a single gas price.

Much of the literature focuses on the case of a monopoly selling into two separate markets with different demand conditions (but identical marginal cost); see Aguirre, Cowan and Vickers (2010) for a recent analysis. Under fairly mild conditions, the resulting uniform price lies between the high and low prices under discrimination. Moreover, price discrimination is usually associated with lower aggregate consumer surplus (across both markets)—although there are exceptions (Cowan, 2012).²⁹

By revealed preference, moving to perfect arbitrage makes a monopolist worse off. The situation is more complex for price-setting oligopoly, and the literature highlights the possibility that price discrimination may reduce industry profits (Corts, 1998). So it is at least conceivable that a shift to a global gas price might be *positive* for LNG exporters (as a group). The impact of price discrimination on social welfare is, in general, ambiguous, and depends, amongst other things, on the fine details of the demand conditions across different markets.³⁰ In the monopoly case, price discrimination is often welfare-reducing—but it is probably more likely to increase welfare under oligopoly. So a move to perfect arbitrage may actually cause global welfare to fall; in any case, it is clear that important distributional effects arise.

However, the assumptions made to obtain these results limit their applicability to LNG markets. First, virtually all of the existing literature focuses on monopoly or price-setting duopoly, neither of which seems a natural choice for LNG mar-

²⁸Stole (2007) provides a useful overview of this literature.

²⁹A smaller number of papers examine third-degree price discrimination by price-setting oligopolies with differentiated products. With symmetric firms, the basic insights from the monopoly case carry over (Holmes, 1989). However, a richer range of outcomes is possible if firms are asymmetric in that they do not rank different markets in the same way, that is, a market is regarded as “strong” by one firm but as “weak” by another firm (Corts, 1998). It is then possible that price discrimination causes prices in both markets to move in the same direction.

³⁰See Bergemann, Brooks and Morris (2013) for a novel welfare analysis for monopoly.

kets. Second, most papers simply assume that firms supply all markets regardless of the degree of price discrimination; this precludes the possibility, for example, that greater price arbitrage might lead to some markets becoming so unattractive to LNG exporters that they are no longer served.³¹ Third, it is typically assumed that each producer has the same marginal cost for each market; this effectively rules out the existence of transport costs, which almost inevitably vary across markets.³² Fourth, particular features of LNG market such as the existence of long-term contracting commitments and its complex supply chain and ownership structure are not modelled. Finally, from a dynamic perspective, the higher profits that firms may derive from the ability to price discriminate can increase their incentives to invest in LNG infrastructure in the first place.

LNG markets seem a fruitful area for research, given their increasing importance and the relative scarcity of existing literature. It would be useful to have more formal results from models of price discrimination with more realistic market structures that can be applied to LNG markets—and elsewhere. It would be particularly interesting to combine economic theory with more extensive market data.

References

- Acharya, Viral V., Lars A. Lochstoer and Tarun Ramadorai (2013). Limits to Arbitrage and Hedging: Evidence from Commodity Markets. *Journal of Financial Economics*, forthcoming.
- Aguirre, Iñaki, Simon Cowan and John Vickers (2010). Monopoly Price Discrimination and Demand Curvature. *American Economic Review* 100, 1601–1615.
- Almoguera, Pedro A., Christopher C. Douglas and Ana María Herrera (2011). Testing for the Cartel in OPEC: Non-Cooperative Collusion or just Non-Cooperative? *Oxford Review of Economic Policy* 27, 144–168.
- Bergemann, Dirk, Benjamin Brooks and Stephen Morris (2013). The Limits of Price Discrimination. Princeton Economic Theory Center Working Paper 052-2013, May 2013.

³¹An exception is Layson (1994) who analyzes market opening in the monopoly case.

³²We conjecture that greater price arbitrage can have important implications for cost efficiency, both in terms of production costs and transport costs. A recent paper by Chen and Schwartz (2013) shows that price discrimination is more likely to raise welfare in the monopoly case when there are marginal-cost asymmetries across markets. Thanks to Simon Cowan for the pointer.

- Borenstein, Severin, James Bushnell, Christopher R. Knittel and Catherine Wolfram (2008). Inefficiencies and Market Power in Financial Arbitrage: A Study of California's Electricity Markets. *Journal of Industrial Economics* 56, 347–378.
- Brito, Dagobert L. and Peter R. Hartley (2007). Expectations and the Evolving World Gas Market. *The Energy Journal* 28, 1–24.
- Chaton, Corinne and Laure Durant-Viel (2013). Real Asset Valuation Under Imperfect Competition: Can We Forget About Market Fundamentals? *Journal of Economics & Management Strategy* 22, 125–139.
- Chen, Yongmin and Marius Schwartz (2013). Differential Pricing When Costs Differ: A Welfare Analysis. Georgetown Economics Working Paper 13-01, April 2013.
- Corts, Kenneth S. (1998). Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment. *RAND Journal of Economics* 29, 306–323.
- Cowan, Simon (2012). Third-Degree Price Discrimination and Consumer Surplus. *Journal of Industrial Economics* 60, 333–345.
- Doane, Michael J. and Daniel F. Spulber (1994). Open Access and the Evolution of the U.S. Spot Market for Natural Gas. *Journal of Law and Economics* 37, 477–517.
- Egging, Ruud, Franziska Holz, Christian von Hirschhausen and Steven Gabriel (2009). Representing GASPEC with the World Gas Model. *The Energy Journal* 2009 Special Issue: World Natural Gas Markets and Trade: A Multi-Modelling Perspective, 97–117.
- GIIGNL (2012). The LNG Industry in 2011. International Group of Liquefied Natural Gas Importers, Paris.
- Holmes, Thomas J. (1989). The Effects of Third-Degree Price Discrimination in Oligopoly. *American Economic Review* 79, 244–250.
- Holz, Franziska, Christian von Hirschhausen and Claudia Kemfert (2008). A Strategic Model of European Gas Supply (GASMOD). *Energy Economics* 30, 766–788.
- Hubert, Franz and Svetlana Ikonnikova (2011). Investment Options and Bargaining Power in the Eurasian Supply Chain for Natural Gas. *Journal of Industrial Economics* 59, 85–116.
- Jensen, James T. (2003). The LNG Revolution. *The Energy Journal* 24, 1–45.

JP Morgan Cazenove (2012). Global LNG. Global Equity Research, London, 13 January 2012.

Layson, Stephen K. (1994). Market Opening under Third-Degree Price Discrimination. *Journal of Industrial Economics* 42, 335–340.

Lise, Wietze and Benjamin Hobbs (2009). A Dynamic Simulation of Market Power in the Liberalized European Natural Gas Market. *The Energy Journal* 2009 Special Issue: World Natural Gas Markets and Trade: A Multi-Modelling Perspective, 119–136.

Nakov, Anton and Galo Nuño (2013). Saudi Arabia and the Oil Market. *The Economic Journal*, forthcoming.

NERA (2012). Macroeconomic Impacts of LNG Exports from the United States. Report for US Department of Energy, Washington DC, 3 December 2012.

Neumann, Anne (2009). Linking Natural Gas Markets – Is LNG Doing Its Job? *The Energy Journal* 2009 Special Issue: World Natural Gas Markets and Trade: A Multi-Modelling Perspective, 187–199.

Platts (2012). LNG Daily, Volume 9, Issue 197, 11 October 2012.

Salant, Stephen W. (1976). Exhaustible Resources and Industrial Structure: A Nash–Cournot Approach to the World Oil Market. *Journal of Political Economy* 84, 1079–1094.

Shleifer, Andrei and Robert Vishny (1997). The Limits of Arbitrage. *Journal of Finance* 52, 35–55.

Silverstovsa, Boriss, Guillaume L'Hégaret, Anne Neumann and Christian von Hirschhausen (2005). International Market Integration for Natural Gas? A Cointegration Analysis of Prices in Europe, North America and Japan. *Energy Economics* 27, 603–615.

Stern, Jonathan P. (editor) (2012). *The Pricing of Internationally Traded Gas*. Oxford: Oxford Institute for Energy Studies.

Stole, Lars (2007). Price Discrimination and Competition. In Mark Armstrong and Robert Porter, *Handbook of Industrial Organization*, Volume 3. Amsterdam: Elsevier, 2221–2299.

Yergin, Daniel and Michael Stoppard (2003). The Next Prize. *Foreign Affairs* 82, 103–114.