Price discrimination and limits to arbitrage: An analysis of global LNG markets

Robert A. Ritz*

Revised version: June 2014

Forthcoming in *Energy Economics*

Abstract

Gas prices around the world vary widely despite being connected by international trade of liquefied natural gas (LNG). Some industry observers argue that major exporters have acted irrationally by not arbitraging prices. This is also difficult to reconcile with a competitive model in which regional price differences exist solely because of transport costs. We show that a model which incorporates market power can rationalize observed prices and trade flows. We highlight how different features of the LNG market limit the ability and/or incentive of other players to engage in arbitrage, including constraints in LNG shipping. We also present some rough estimates of market power in short-term sales by Qatar (to Japan and the UK, respectively), and discuss the potential impact of US LNG exports.

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

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

*Affiliation: Faculty of Economics & Energy Policy Research Group, University of Cambridge, Sidgwick Avenue, Cambridge CB3 9DD, UK; E-mail: rar36@cam.ac.uk; Telephone: +44 (0)1223 335289

1 Introduction

Motivation for the analysis. Not so long ago, there was a widespread conjecture that gas prices around the world would converge. The idea was that sea-based trade in liquefied natural gas (LNG) would connect regional markets—notably 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 an increase in demand for natural gas. Global investment in LNG infrastructure—liquefaction and regasification capacity—has risen, and such facilities are now spread across more countries. The size of the 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 have varied widely; significant price differentials have existed for years, and have become more pronounced since the Fukushima accident of March 2011 (IGU, 2013). The 2012 average natural gas price was roughly US$16/MMBtu in Japan, $9 in Europe but only $3 in the US. Some expect large price disparities to persist, including the International Energy Agency in modelling scenarios for 2020 and 2035. In short, the gas market appears far from global.

For the case of the US, the reasons for price divergence are quite clear. First, the

[1] Useful 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.

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

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]

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 = |price differential| × 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. (These deliveries are estimated as 75% higher than contractual obligations.) LNG producers have, apparently, been 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 observed over the last years. Consider again the case of Qatar, the largest LNG producer with a global market share of around 30%, and its sales to Japan and UK, the two largest routes of short-term LNG, with global shares of 10% and 6% respectively (GIIGNL, 2012). Figure 1 shows the differential between the Platts JKM (Japan Korea Marker) spot price for Asian LNG and the UK NBP (National

[3] Several LNG liquefaction facilities are currently under development or pending US regulatory/political approval.

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

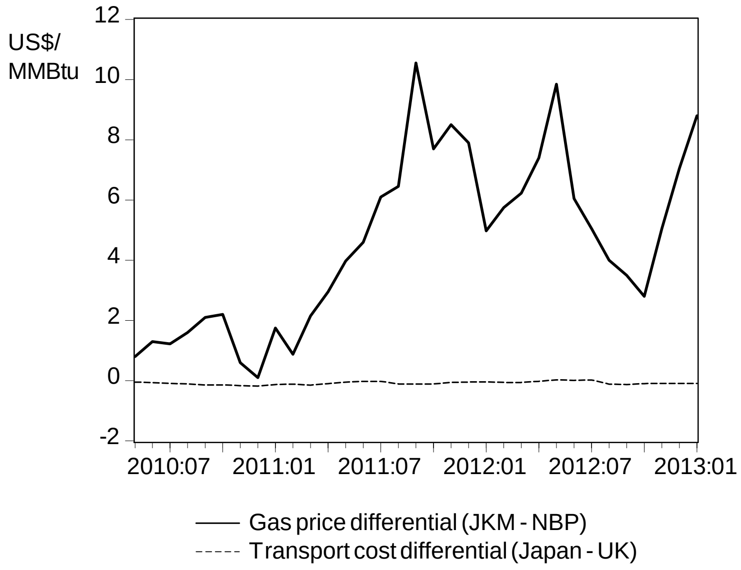

Balancing Point) hub price, plotted against the difference in transport costs between shipping from Qatar to Japan and Qatar to the UK, respectively.[4,5] Prices are up to $10/MMBtu higher in Asia than in the UK, while the corresponding transport costs are approximately identical. Perfect competition, by contrast, predicts that these two differentials should coincide.[6]

The perfectly competitive model cannot account for other producers' behaviour. For example, Peru is in a similar position to Qatar in that its Europe-Asia transport-cost differentials are usually very small (Platts, 2012), and yet it makes short-term

[4]Platts JKM reflects the spot price of LNG in Japan and South Korea, reported as a single market-wide number. In practice, there may be some variation in prices for different buyer-seller pairings, but information on such individual transactions is generally unavailable. Also widely reported is an LNG price based on the Japanese Crude Cocktail (JCC); this reflects oil-linked pricing formulae that underlie long-term supply contracts—rather than spot trades for uncommitted LNG on which this paper focuses. See Stern (2012) for a detailed overview of how gas-pricing mechanisms vary across regions, including on Asian gas price formation.

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

[6]This theoretical prediction should 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.

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.

sales to both markets (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.

Overview of the results. In this paper, we instead suggest that regional price differentials 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, 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.

The 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; this may optimally leave prices across markets far apart. This argument extends straightforwardly to more than two export markets and to capacity constraints on production; moreover, it does not depend on any particular mode of strategic competition (such as Cournot-Nash competition).[7]

We provide a simple new formula for relative prices between any two export markets, in terms of a seller's transport costs and price elasticities of demand. It deals with a key challenge—the shortage of available data on LNG markets. In particular, our basic approach does not rely on production cost data, or on any assessment of whether or not an exporter is capacity-constrained. Also, while our exposition focuses on gas and LNG, it will become clear that our approach and results could also find application in other contexts, including other "commodity" markets as well as settings with differentiated products.

Market power can thus rationalize observed gas prices and trade flows by tracing

[7] We do not suggest that LNG producers are colluding, but rather that at least some of them 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".

them to local 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". (Local demand conditions play no role in the competitive model, in which price differences are *solely* driven by transport costs.) Our rough empirical estimates for Qatar yield Lerner indices (price-cost mark-ups) of around 63% for short-term LNG sales to Japan compared to 33% for sales to the UK (in 2012). We discuss in detail 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.

We can 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, an influential 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.[8]

Our model of price discrimination with imperfect competition is static and thus does not explicitly capture intertemporal features of the gas market (such as storage).[9] In practice, there may be a dynamic interaction between short-term 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 producers may have been reluctant to push down short-term prices in Asia insofar as this would make it more difficult to sustain

[8] NERA (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.

[9] See Chaton and Durant-Viel (2013) for a recent analysis of the value of gas storage when firms have market power.

"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 convergence in natural gas markets, though few of these focus specifically on LNG. Silivverstovsa 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.[10] Others have focused on price convergence within regional markets; for example, Doane and Spulber (1994) employ similar techniques to find an integrated, national market for natural gas in the US. By contrast, we offer an economic-theory perspective on price non-convergence, including a view to explaining observed prices since March 2011.

There 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). Newbery (2008) examines the impact of carbon pricing on market power in the EU gas market (but does not consider LNG trade). In related work, Growitsch, Hecking and Panke (2013) analyze the impact of disruptions in the LNG supply chain at the Strait of Hormuz on regional market outcomes, and also emphasize market power in the Japanese market. We also employ 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.

Finally, there is related literature on market power in crude oil which estimates

[10] There does not appear to be any econometric analysis of gas-price convergence post-2011. We conjecture that international price correlations have declined significantly (compared to the mid/late 2000s), and that the Fukushima accident represents a structural break.

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). 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 price arbitrage in LNG.

This paper proceeds as follows. Section 2 presents a model of pricing across different export markets, and gives some empirical estimates for the case of Qatar. Section 3 provides a detailed discussion of limits to arbitrage in global LNG markets. Section 4 offers concluding remarks on the potential effects of more LNG arbitrage, in the future, on prices, industry profits, and social welfare—and thereby highlights several shortcomings of the existing literature on third-degree price discrimination.

2 A profit-maximizing LNG exporter

Setup of the model. Consider an LNG producer selling output into $M \geq 2$ export markets. Let $p_\ell^k(x_\ell^k, y_\ell^k, X_\ell^{-k}, Y_\ell^{-k}; \boldsymbol{\theta}_\ell)$ denote the inverse demand function producer k faces in market ℓ ($\ell = 1, .., M$), where p_ℓ^k is the "spot" price of LNG, x_ℓ^k is the quantity sold by the producer in the short-term market, while y_ℓ^k is the quantity the producer has pre-committed to sell in market ℓ by way of long-term contracts (which we here take as given). Analogously, X_ℓ^{-k} is the vector of outputs sold by other producers in the spot market while Y_ℓ^{-k} captures their long-term commitments.[11] Other factors that affect demand in market ℓ are summarized by the vector $\boldsymbol{\theta}_\ell$. For example, this might include prices of pipeline-based gas, coal, oil

[11] From a buyer's point of view, LNG from short-term markets and long-term contracts may be imperfect substitutes.

and other substitutes, the business cycle, other demand shocks, and the weather.[12]

In the interest of generality, the model allows—but does not require—prices to be producer-specific, that is, for there to be more than a single price in a market. Most existing analyses of natural gas markets assume a single price in each market, and all our modelling results hold for such situations (by setting $p_\ell^k = p_\ell$). Our model can also apply to situations in which there is price variation within a market, e.g., because of quality differentials between the products of various sellers, which could be due to differences in the energy content of delivered gas.

The producer's cost function $C^k(\sum_{\ell=1}^{M}(x_\ell^k + y_\ell^k))$ depends on the sum of total quantities sold in all M export markets, including spot market sales as well as long-term commitments. Production may be subject to a capacity constraint such that $\sum_{\ell=1}^{M}(x_\ell^k + y_\ell^k) \leq \overline{Q}^k$.[13] In addition to this, producer k incurs a market-specific transport cost τ_ℓ^k per unit of output sold to market ℓ.[14] 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).

Our analysis here assumes that the LNG producer has access to shipping at the respective transport rates; this is an appropriate assumption for major exporters such as Qatar. We return to the question of access to the LNG tanker market in our discussion of limits to price arbitrage by other parties in Section 3.

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:

$$\max_{\{x_\ell^k\}_{\ell=1}^{M}} \Pi^k = \sum_{\ell=1}^{M} p_\ell^k x_\ell^k - C^k(\sum_{\ell=1}^{M}(x_\ell^k + y_\ell^k)) - \sum_{\ell=1}^{M} \tau_\ell^k x_\ell^k$$

$$\text{subject to } \sum_{\ell=1}^{M}(x_\ell^k + y_\ell^k) \leq \overline{Q}^k.$$

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

[13]Adding a production constraint in terms of minimum throughput to ensure smooth operation ("must run"), $\sum_{\ell=1}^{M}(x_\ell^k + y_\ell^k) \geq \underline{Q}^k$, would not affect the following results.

[14]The results would remain unchanged with an affine transport cost function of the form $T_\ell + \tau_\ell^k x_\ell^k$, where T_ℓ represents a fixed cost.

We assume, without much loss of additional economic insight, that this problem is well-behaved with an interior solution for each of the M export markets.[15] The Langrangean for constrained optimization can be written as $L^k = \Pi^k + \lambda^k(\overline{Q}^k - \sum_{\ell=1}^{M}(x_\ell^k + y_\ell^k))$, 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.

Profit-maximization analysis. The optimal output choice $\widehat{x}_\ell^k$ by producer k in market ℓ satisfies the first-order condition $MR_\ell^k - MC^k - \tau_\ell^k - \lambda^k = 0$, where MR_ℓ^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

$$MR_i^k - \tau_i^k = MR_j^k - \tau_j^k.$$

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, 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 ℓ can be written as $MR_\ell^k = p_\ell^k \left(1 - 1/\eta_\ell^k\right)$, where η_ℓ^k is the own-price elasticity of producer k's demand (alternatively, the elasticity of its residual demand). From here on, this elasticity is understood to be evaluated at producer k's optimally chosen output $\widehat{x}_\ell^k$, as well as at the *actual* levels of short-term

[15] We are interested in explaining prices associated with empirically observed non-zero trade flows to particular markets—rather than working out which markets a producer should in fact be selling to, and which ones not. For expositional purposes, we therefore focus on the $M \geq 2$ markets to which positive sales occur, rather than the general problem where the seller may choose to sell to only M markets of a total number of $M' > M$ possible markets. In any case, our fundamental conditions for profit-maximization also apply in the general problem to any markets which the producer does in fact choose to sell to.

output of other producers and corresponding actual long-term commitments, that is, $\eta_\ell^k = \eta_\ell^k(\widehat{x}_\ell^k, y_\ell^k, X_\ell^{-k}, Y_\ell^{-k}; \boldsymbol{\theta}_\ell) > 1$.[16]

Proposition 1 *A profit-maximizing producer k sells into $M \geq 2$ export markets with a common marginal cost (and possibly subject to a capacity constraint):*
(A) In any two markets i and j, profit-maximizing prices (p_i^k, p_j^k) satisfy

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

where (τ_i^k, τ_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 (τ_i^k, τ_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).

Proof. For (A), rewrite the fundamental equation for profit-maximization $MR_i^k - \tau_i^k = MR_j^k - \tau_j^k$ using the expression $MR_l^k = p_l^k\left(1 - 1/\eta_\ell^k\right)$ for $l = i, j$ to obtain

$$\begin{aligned}
& p_i^k \frac{(\eta_i^k - 1)}{\eta_i^k} - \tau_i^k = p_j^k \frac{(\eta_j^k - 1)}{\eta_j^k} - \tau_j^k \\
\Longrightarrow \quad & (p_i^k - p_j^k)\frac{(\eta_i^k - 1)}{\eta_i^k} = p_j^k\left[\frac{(\eta_j^k - 1)}{\eta_j^k} - \frac{(\eta_i^k - 1)}{\eta_i^k}\right] + (\tau_i^k - \tau_j^k) \\
\Longrightarrow \quad & (p_i^k - p_j^k)\frac{(\eta_i^k - 1)}{\eta_i^k} = p_j^k\left(\frac{1}{\eta_i^k} - \frac{1}{\eta_j^k}\right) + (\tau_i^k - \tau_j^k),
\end{aligned}$$

from which the result for the relative price $(p_i^k - p_j^k)/p_j^k$ follows immediately.
For (B), observe that k's first-order condition for market i, $MR_i^k - MC^k - \tau_i^k - \lambda^k = 0$, can be rewritten, again using $MR_i^k = p_i^k\left(1 - 1/\eta_\ell^k\right)$, to give

$$\begin{aligned}
& p_i^k = \frac{\eta_i^k}{(\eta_i^k - 1)}\left[\tau_i^k + \left(MC^k + \lambda^k\right)\right] \\
\Longrightarrow \quad & (p_i^k - \tau_i^k) = \frac{1}{(\eta_i^k - 1)}\left[\tau_i^k + \eta_i^k\left(MC^k + \lambda^k\right)\right],
\end{aligned}$$

[16] A necessary condition for profit-maximization is that producer demand remains price-elastic in each market, $\eta_\ell^k > 1$. (Otherwise the producer could profitably reduce output.) Market-level demand elasticities can be significantly lower; see also our discussion near the end of this section.

which decreases in the elasticity η_i^k, and hence satisfies $(p_i^k - \tau_i^k) \in \left[\left(MC^k + \lambda^k\right), \infty\right)$ given the limiting cases, respectively, of $\eta_i^k \to \infty$ and $\eta_i^k \to 1$. So any price and transport cost satisfying $p_i^k \geq \tau_i^k$ can be rationalized for some marginal production (shadow) cost $\left(MC^k + \lambda^k\right) \geq 0$ and an appropriately chosen elasticity η_i^k. The same argument applies for market j. ■

Understanding profit-maximizing prices. The formula for relative prices from (A) is rather general: it does not rely on any specific functional-form assumptions 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 differ across export markets.[17] An informational advantage is that it does not feature production costs. (The result also does 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 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 $\eta_\ell^k \to \infty$ in each export market $\ell = 1, .., M$. 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 $\left(p_i^k - p_j^k\right) = \left(\tau_i^k - \tau_j^k\right)$, and netbacks are equalized.

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

[17] Although we do not emphasize this aspect here, the same result would also apply to the two-stage setting of Allaz and Vila (1993), where firms, in the first stage, can commit to selling output in forward markets, and, then, in the second stage, choose their production levels.

Third, assume that transport costs to two markets are identical, $\tau_i^k = \tau_j^k = \widehat{\tau}$. Relative price then satisfies $\left(p_i^k - p_j^k\right)/p_j^k = \left(\eta_j^k - \eta_i^k\right)/\left(\eta_i^k - 1\right)\eta_j^k$ 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 $\tau_i^k \geq \tau_j^k$, 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 $\text{sign}\left(p_i^k - p_j^k\right) = \text{sign}\left(\tau_i^k - \tau_j^k\right)$, and can also feature "excess volatility" in prices, by going beyond the implication from the competitive model that $\text{var}\left(p_i^k - p_j^k\right) = \text{var}\left(\tau_i^k - \tau_j^k\right)$.

While the present paper focuses on global gas pricing, the result from Proposition 1 could also find application in other settings, both "commodity" markets and those in which firms' products are significantly differentiated.

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 ℓ can generate any non-negative price-cost margin, ranging from zero (when $\eta_\ell^k \to \infty$) to arbitrarily large (when $\eta_\ell^k \to 1$), regardless of the underlying details of the producer's costs. (A negative price-cost margin cannot be profit-maximizing in the present model.) 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 LNG sales to Asia and Northwest Europe

(i = Asia, j = NW Europe). Let prices $p_i^k = 16$ and $p_j^k = 9$, and assume, for simplicity, that transport costs for producer k are identical, $\tau_i^k = \tau_j^k = \widehat{\tau}$. Such prices can thus be explained by the producer having relatively greater market power in Asia, that is, $\eta_i^k < \eta_j^k$. Indeed, it is not difficult to check that these relative prices can be rationalized by a pair of elasticities $\eta_i^k = 2$ and $\eta_j^k = 9$. (Note that the precise value of $\widehat{\tau}$ does not matter for this to hold.) Clearly, similar examples can also be constructed to explain the behaviour of other producers in other markets, including for cases where $\tau_i^k \neq \tau_j^k$.

Note that the above move from the modelling results to real-world data implies an assumption that there is a single price *within* each regional market, i.e., each seller receives the same price. In particular, our price data in form of the Platts JKM price as well as the UK NBP price come as a single number for each of the two respective markets (at each point in time). Such single prices do not capture any price variation across different buyer-seller pairings, so a "law of one price" holds inside each market. Recall, however, that that our theoretical argument does not hinge upon this.[18]

Thus recent claims that major LNG exporters 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 Europe to keep prices in Asia high.[19]

Why might producers have greater market power in Asia, notably 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

[18]In this sense, our empirical analysis is more restrictive than our theoretical model. The assumption of single price in each regional market seems a good assumption for hub-based markets, for example, in Northwest Europe, but probably more restrictive for less liquid markets in Asia. Unfortunately, we do not have access to data on individual transactions.

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

fairly general conditions, an upward shift in market demand (captured formally by a change in $\boldsymbol{\theta}_i$) 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.[20]

To illustrate, consider a Cournot-Nash model of export market i in which $N^i \geq 2$ firms face a common demand curve $p_i(X_i;\theta_i) = \theta_i + f_i(X_i)$, where $X_i \equiv \sum_k x_i^k$, and θ_i shifts consumers' willingness-to-pay, and suppose producer k has constant marginal cost MC^k.[21] The first-order condition for k is $MR_i^k - MC^k - \tau_i^k = 0$, where $MR_i^k = [\theta_i + f_i(X_i)] + f_i'(X_i)\widehat{x}_i^k < p_i$. Summing over all firms and rearranging gives $N^i[\theta_i + f_i(\widehat{X}_i)] + f_i'(\widehat{X}_i)\widehat{X}_i = \sum_k MC^k + \sum_k \tau_i^k$. Differentiating with respect to a positive demand shock $d\theta_i > 0$ shows that equilibrium demand/imports rise since

$$\frac{d\widehat{X}_i}{d\theta_i} = \frac{N^i}{\left[-f_i'(\widehat{X}_i)\right]\left[(N^i+1) + \widehat{X}_i f_i''(\widehat{X}_i)/f_i'(\widehat{X}_i)\right]} > 0,$$

where the denominator is positive by the stability of equilibrium. The market price responds according to $dp_i/d\theta_i = \partial p_i/\partial\theta_i + (\partial p_i/\partial X_i)|_{X_i=\widehat{X}_i}\,(d\widehat{X}_i/d\theta_i)$, so, using $\partial p_i/\partial\theta_i = 1$ and $\partial p_i/\partial X_i = f_i'(X_i)$, we obtain

$$\frac{dp_i}{d\theta_i} = \frac{1 + \widehat{X}_i f_i''(\widehat{X}_i)/f_i'(\widehat{X}_i)}{\left[(N^i+1) + \widehat{X}_i f_i''(\widehat{X}_i)/f_i'(\widehat{X}_i)\right]}.$$

So the condition for price to rise—and producer-specific elasticities to fall—is:

$$\frac{dp_i}{d\theta_i} > 0 \Longleftrightarrow \frac{d\eta_i^k}{d\theta_i} < 0 \text{ (for all } k) \iff \frac{\widehat{X}_i f_i''(\widehat{X}_i)}{f_i'(\widehat{X}_i)} > -1.$$

[20] Another potential explanation—albeit one rejected by the available data—is that capacity constraints in an importing country's regasification infrastructure drive price above marginal cost. This *could* explain regional price divergences even in the absence of market power. At a global level, however, the rate of capacity utilisation of regasification terminals has been stable, and rather low, at around 40% since 2000. In Japan, annual LNG imports increased by 24% between 2010 and 2012 without significant capacity additions; indeed utilisation rates rose only moderately up to 49%. Utilisation is higher for some importing countries, notably China and India, but at a similar or lower level for other such as Argentina, Korea and the UK. (All based on data in IGU, 2013.)

[21] For simplicity, this illustration assumes away long-term contracts and third-party arbitrage with other markets (although the points made would apply more generally).

The latter condition requires that demand is "not too convex" (around equilibrium) and is satisfied, for example, by any linear or concave demand curve. More formally, it corresponds to "log-concavity" of demand, that is, the direct demand curve is such that $\log D_i(p_i)$ is concave in price, which is a common assumption in economic theory. In this sense, higher demand exacerbates sellers' market power. By contrast, in a competitive market $dp_i/d\theta_i = 0$, given constant returns to scale.[22]

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, lower elasticities. Furthermore, Asian buyers have fewer possibilities to substitute for LNG, notably because of more limited access to Russian pipeline gas.

Estimating producer-specific demand elasticities. A feature of the above model is that the pair of elasticities to rationalize the data is, in general, not unique. In the above numerical example, the data $p_i^k = 16$, $p_j^k = 9$, and $\tau_i^k = \tau_j^k = \widehat{\tau}$ 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; setting $(\eta_i^k, \eta_j^k) = (1\frac{1}{2}, 2\frac{5}{11})$ is another. Loosely put, getting the relative elasticities across markets correct matters more than their absolute values. Note, however, that certain bounds on elasticities must be satisfied. To see this, let $\tau_i^k = \tau_j^k = \widehat{\tau}$, and, suppose, without loss of generality that $p_i^k > p_j^k$. To be able to rationalize these data, the producer-specific elasticity for market i must be sufficiently low in that $\eta_i^k < p_i^k/(p_i^k - p_j^k)$.

Pinning down unique values for the elasticities requires more information and/or additional modelling assumptions. We here discuss two approaches, which may be substitutes or complements, and give some rough estimates of Qatar's market power.

First, recall that the above results did not rely on production cost informa-

[22] For competitive markets with (strictly) upward-sloping supply curves, we would have $dp_i/d\theta_i > 0$. But then the price/demand increase in market i would, in equilibrium, increase marginal cost for *all* other markets, such that $dp_j/d\theta_i = dp_i/d\theta_i > 0$ (for all $j \neq i$). So equilibrium price differentials $(p_i - p_j) = (\tau_i - \tau_j)$ between any two markets remain unchanged.

tion, 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 ℓ, using the relationship $MR_\ell^k = p_\ell^k \left(1 - 1/\eta_\ell^k\right)$, to obtain $\eta_i^k = p_i^k / \left[(p_i^k - \tau_i^k) - (MC^k + \lambda^k)\right]$. In addition to prices (p_i^k, p_j^k) and transport costs (τ_i^k, τ_j^k), knowledge of $(MC^k + \lambda^k)$ would identify each of the elasticities, η_i^k and η_j^k, and yield a unique way of rationalizing the data. In the above numerical example, for instance, knowledge that $MC^k + \lambda^k = 6$ would select $(\eta_i^k, \eta_j^k) = (2, 9)$ as the unique solution. Perhaps most realistically, an assessment that the producer is not capacity-constrained, i.e., $\lambda^k = 0$, together with data on $(p_\ell^k, \tau_\ell^k, MC^k)$ identifies η_ℓ^k for each individual market.[23]

We here present some illustrative results for the case of Qatari exports to Japan and the UK during 2012 ($p_i^k = 16$, $p_j^k = 9$). These are based on indicative cost estimates in IEA (2009), which give the combined unit cost of production, liquefaction, and regasification of LNG at around \$3.00/MMBtu in 2008 US\$.[24,25] We here assume that Qatar is not capacity-constrained ($\lambda^k = 0$), and take this estimate as a proxy for its marginal cost. To account for possible cost inflation between 2008 and 2012, we employ $MC^k = 3.90$ (corresponding to year-on-year cost increases of just under 7%). For transport costs, we rely on the data used above and set $\tau_i^k = 2.10$ for Japan, and $\tau_j^k = 2.15$ for the UK market (based on 2012 averages). Taken together, this information yields estimates of the producer-specific elasticities $\eta_i^k = 1\frac{3}{5}$ and $\eta_j^k = 3\frac{3}{59}$. These can be equivalently expressed as Lerner indices (price-cost margins) of $L_i^k \equiv \left[(p_i^k - \tau_i^k) - MC^k\right]/p_i^k = (\eta_i^k)^{-1} \simeq 63\%$ for Japan and $L_j^k = (\eta_j^k)^{-1} \simeq 33\%$ for the UK. We emphasize that these market-power estimates

[23] Put differently, for each seller, there are M first-order conditions but $M+2$ unknowns, namely, $\{\eta_\ell^k\}_{\ell=1}^M$, 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 $M+1$ unknowns.

[24] IEA (2009) is the only publicly available source of (fairly) recent (and only "indicative") LNG cost data we have been able to access.

[25] The unit cost of production for Qatar is estimated to be close to zero, as project costs are taken to be covered by the output of gas condensate and other liquids, so the large bulk of costs stems from liquefaction and regasification.

are only indicative; however, even accounting for potential shortcomings in the assumptions and data, it seems clear that (i) price-cost margins are significant for both markets, and (ii) margins are considerably higher in Japan than for the UK.

Second, 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 $\eta_\ell^k = \eta_\ell / s_\ell^k$, where η_ℓ is the price elasticity of *market* demand and $s_\ell^k \in (0,1)$ is the producer's market share—in the appropriately defined market ℓ. Conditional on settling the issue of market definition, data on market shares are generally easier to obtain, and it is usually easier to estimate a market-level elasticity than a producer-specific elasticity (or a set of producer elasticities).

This second approach can act as a substitute for the previous production-cost approach, namely as an alternative way of pinning down the precise value of η_ℓ^k for each seller in each market. But it can also be used in conjunction with the first approach to test the validity of particular models of competition. More specifically, it could be used to (i) provide a direct estimate, even in the absence of any data on prices or costs, of η_ℓ^k for a producer-market pairing; (ii) select a preferred pair from the set of producer-specific elasticities (η_i^k, η_j^k) obtained using our basic model which utilizes data on prices and transport costs; (iii) test whether a particular model of competition, say Cournot-Nash, is consistent with the values of (η_i^k, η_j^k) obtained by our first approach using production-cost information.

To illustrate using our numerical example from above, suppose it is estimated that two markets have identical overall price elasticities $\eta_i = \eta_j = \frac{1}{2}$ (say for natural gas, obtained both via LNG and pipelines where available); 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 $(\eta_i, \eta_j) = (\frac{1}{5}, \frac{9}{10})$.

To summarize, we have presented a basic model that requires only minimal information on prices and transport costs. This model can be augmented using additional input data on production costs and/or competitive conditions.[26] We have derived some illustrative estimates of Qatar's short-term market power in Japan and the UK. It is clear that the general methodology could easily be applied to other producers and trade routes insofar as sufficient data were available to the analyst. We return to the impact of price discrimination under different competitive conditions in our concluding remarks below.

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 producers still normally restrict the resale of their exports in a way that prevents "secondary" trading 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.[27] Al-

[26] It is also possible, in principle, to estimate the elasticity for an individual markets, η_ℓ^k, or the entire vector of elasticities $\{\eta_\ell^k\}_{\ell=1}^M$, using econometric techniques. This would require time-series data on prices and quantities in each market of interest (probably ideally panel data), as well as control variables (probably including data on producers' long-term contractual commitments).

[27] Our model above assumed 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.

though 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.[28] 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.[29] Commissioning a new LNG carrier is a costly and uncertain undertaking, with a significant time lag due to construction ("time to build").

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.[30] Second, ownership arrangements along the LNG supply chain are much more complex than in any 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

Moreover, our numerical illustrations use publicly available shipping rates as being representative.

[28]In some cases, there may also be issues of compatibility: The existing global fleet of LNG tankers has widely varying vintages, with some vessels dating from the 1970s and some very new tankers; vessels come in different shapes and sizes, and not all LNG import terminals are able to receive deliveries from all types of LNG tanker.

[29]There is also a potential feedback effect: Short-term LNG typically involves longer distances than trade from long-term contracts, so 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.

[30]We do not have data on downstream margins of integrated LNG players.

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, can 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; Acharya, Lochstoer and Ramadorai, 2013).

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.[31] For example, the optimal arbitrage strategy $\widehat{\alpha} \equiv \arg\max_{\alpha} \{\alpha [p_i(-\alpha) - p_j(\alpha)]\}$ of buying α units in low-price market j to sell in high-price market i leaves $p_i(-\widehat{\alpha}) \neq p_j(\widehat{\alpha})$ 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 discussion

Summary of the analysis. 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 years, becoming even more pronounced since the Fukushima accident of March 2011. Some industry observers have 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 has presented the first attempt in the literature to address this puzzle. It shows that observed prices and trade flows can be explained by LNG exporters' market power; indeed, it is quite difficult to see how the data could be rationalized

[31] See also Borenstein, Bushnell, Knittel and Wolfram (2008) on pricing in Californian electricity markets around the time of the Enron collapse; they show that arbitrage oppurtunities 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.

without incorporating 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. Rough estimates for Qatar suggest significant market power, albeit at different levels, for the Japanese and UK markets. We have argued that, in addition, a combination of incentives, market power, and other constraints tends to work against international price arbitrage by LNG buyers and third-party traders.

Discussion of LNG market developments, and future research. 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, as well as from Russian export capacity. Yet other LNG projects, notably in Australia, have higher-than-projected costs which may dampen future 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 in Japan to introduce LNG futures contracts would facilitate hedging and arbitrage.

A natural question therefore is, how will greater LNG arbitrage affect the global gas market? The existing theoretical literature on third-degree price discrimination offers some partial answers.[32] 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

[32] Stole (2007) provides a useful overview of this literature.

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).[33]

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.[34] 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 gas markets. First, virtually all of the existing literature focuses on monopoly or price-setting duopoly, neither of which seems a natural choice for LNG markets. 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 exporters that they are no longer served.[35] Third, it is typically assumed that each

[33] A 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.

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

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

producer has the same marginal cost for each market; this effectively rules out the existence of transport costs, which almost inevitably vary across markets.[36] 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.

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

Acknowledgements: Thanks are due to Richard Tol (the Editor) and two referees, Alberto Behar, Chi Kong Chyong, Simon Cowan, Meredith Crowley, Philipp Koenig, Yassine Lefouili, 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 seminar participants at Cambridge, EWI Cologne, the 2013 EPRG Winter Seminar, and the 2014 Toulouse Energy Conference for their feedback. All views expressed and any errors are my own.

References

Acharya, Viral V., Lars A. Lochstoer and Tarun Ramadorai (2013). Limits to Arbitrage and Hedging: Evidence from Commodity Markets. *Journal of Financial*

[36] We 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.

Economics 109, 441–465.

Aguirre, Iñaki, Simon Cowan and John Vickers (2010). Monopoly Price Discrimination and Demand Curvature. *American Economic Review* 100, 1601–1615.

Allaz, Blaise and Jean-Luc Vila (1993). Cournot Competition, Forward Markets and Efficiency. *Journal of Economic Theory* 59, 1–16.

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.

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.

Growitsch, Christian, Harald Hecking and Timo Panke (2013). Supply Disruptions and Regional Price Effects in a Spatial Oligopoly – An Application to the Global Gas Market. EWI–WP 2013/08, *Review of International Economics*, forthcoming

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.

IGU (2013). World LNG Report – 2013 Edition. International Gas Union, Norway.

International Energy Agency (2009). *World Energy Outlook*. Paris: OECD/IEA.

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* 123, 1333–1362.

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.

Newbery, David M. (2008). Climate Change Policy and its Effect on Market Power in the Gas Market. *Journal of the European Economic Association* 6, 727–751.

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.

Siliverstovsa, 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.