

Modeling and Valuation of Energy Structures

Daniel Mahoney

Analytics, Econometrics, and Numerics



Modeling and Valuation of Energy Structures

Applied Quantitative Finance series

Applied Quantitative Finance is a new series developed to bring readers the very latest market tested tools, techniques and developments in quantitative finance. Written for practitioners who need to understand how things work “on the floor”, the series will deliver the most cutting-edge applications in areas such as asset pricing, risk management and financial derivatives. Although written with practitioners in mind, this series will also appeal to researchers and students who want to see how quantitative finance is applied in practice.

Also available

Oliver Brockhaus

EQUITY DERIVATIVES AND HYBRIDS

Markets, Models and Methods

Enrico Edoli, Stefano Fiorenzani and Tiziano Vargiolu

OPTIMIZATION METHODS FOR GAS AND POWER MARKETS

Theory and Cases

Roland Lichters, Roland Stamm and Donal Gallagher

MODERN DERIVATIVES PRICING AND CREDIT EXPOSURE ANALYSIS

Theory and Practice of CSA and XVA Pricing, Exposure Simulation and Backtesting

Zareer Dadachanji

FX BARRIER OPTIONS

A Comprehensive Guide for Industry Quants

Ignacio Ruiz

XVA DESKS: A NEW ERA FOR RISK MANAGEMENT

Understanding, Building and Managing Counterparty and Funding Risk

Christian Crispoldi, Peter Larkin and Gérald Wigger

SABR AND SABR LIBOR MARKET MODEL IN PRACTICE

With Examples Implemented in Python

Adil Reghai

QUANTITATIVE FINANCE

Back to Basic Principles

Chris Kenyon and Roland Stamm

DISCOUNTING, LIBOR, CVA AND FUNDING

Interest Rate and Credit Pricing

Marc Henrard

INTEREST RATE MODELLING IN THE MULTI-CURVE FRAMEWORK

Foundations, Evolution and Implementation


Modeling and Valuation of Energy Structures

Analytics, Econometrics, and Numerics

Daniel Mahoney

Director of Quantitative Analysis, Citigroup, USA

palgrave
macmillan



© Daniel Mahoney 2016

Softcover reprint of the hardcover 1st edition 2016 978-1-137-56014-8

All rights reserved. No reproduction, copy or transmission of this publication may be made without written permission.

No portion of this publication may be reproduced, copied or transmitted save with written permission or in accordance with the provisions of the Copyright, Designs and Patents Act 1988, or under the terms of any licence permitting limited copying issued by the Copyright Licensing Agency, Saffron House, 6–10 Kirby Street, London EC1N 8TS.

Any person who does any unauthorized act in relation to this publication may be liable to criminal prosecution and civil claims for damages.

The author has asserted his right to be identified as the author of this work in accordance with the Copyright, Designs and Patents Act 1988.

First published 2016 by
PALGRAVE MACMILLAN

Palgrave Macmillan in the UK is an imprint of Macmillan Publishers Limited, registered in England, company number 785998, of Houndmills, Basingstoke, Hampshire RG21 6XS.

Palgrave Macmillan in the US is a division of St Martin's Press LLC, 175 Fifth Avenue, New York, NY 10010.

Palgrave Macmillan is the global academic imprint of the above companies and has companies and representatives throughout the world.

Palgrave® and Macmillan® are registered trademarks in the United States, the United Kingdom, Europe and other countries.

ISBN 978-1-349-56688-4 ISBN 978-1-137-56015-5 (eBook)
DOI 10.1057/9781137560155

This book is printed on paper suitable for recycling and made from fully managed and sustained forest sources. Logging, pulping and manufacturing processes are expected to conform to the environmental regulations of the country of origin.

A catalogue record for this book is available from the British Library.

A catalog record for this book is available from the Library of Congress.

To Cathy, Maddie, and Jack

This page intentionally left blank

Contents

<i>List of Figures</i>	xi
<i>List of Tables</i>	xiii
<i>Preface</i>	xiv
<i>Acknowledgments</i>	xviii
1 Synopsis of Selected Energy Markets and Structures	1
1.1 Challenges of modeling in energy markets	1
1.1.1 High volatilities/jumps	1
1.1.2 Small samples	2
1.1.3 Structural change	3
1.1.4 Physical/operational constraints	4
1.2 Characteristic structured products	4
1.2.1 Tolling arrangements	4
1.2.2 Gas transport	6
1.2.3 Gas storage	7
1.2.4 Load serving	9
1.3 Prelude to robust valuation	11
2 Data Analysis and Statistical Issues	12
2.1 Stationary vs. non-stationary processes	12
2.1.1 Concepts	12
2.1.2 Basic discrete time models: AR and VAR	22
2.2 Variance scaling laws and volatility accumulation	29
2.2.1 The role of fundamentals and exogenous drivers	31
2.2.2 Time scales and robust estimation	33
2.2.3 Jumps and estimation issues	34
2.2.4 Spot prices	39
2.2.5 Forward prices	42
2.2.6 Demand side: temperature	43
2.2.7 Supply side: heat rates, spreads, and production structure	46
2.3 A recap	47
3 Valuation, Portfolios, and Optimization	48
3.1 Optionality, hedging, and valuation	48
3.1.1 Valuation as a portfolio construction problem	48
3.1.2 Black Scholes as a paradigm	52

3.1.3	Static vs. dynamic strategies	58
3.1.4	More on dynamic hedging: rolling intrinsic	68
3.1.5	Market resolution and liquidity	75
3.1.6	Hedging miscellany: greeks, hedge costs, and discounting	79
3.2	Incomplete markets and the minimal martingale measure	85
3.2.1	Valuation and dynamic strategies	86
3.2.2	Residual risk and portfolio analysis	88
3.3	Stochastic optimization	101
3.3.1	Stochastic dynamic programming and HJB	101
3.3.2	Martingale duality	106
3.4	Appendix	111
3.4.1	Vega hedging and value drivers	111
3.4.2	Value drivers and information conditioning	113
4	Selected Case Studies	118
4.1	Storage	118
4.2	Tolling	121
4.3	Tolling	128
4.3.1	(Monthly) Spread option representation of storage	128
4.3.2	Lower-bound tolling payoffs	129
5	Analytical Techniques	131
5.1	Change of measure techniques	131
5.1.1	Review/main ideas	131
5.1.2	Dimension reduction/computation facilitation/estimation robustness	135
5.1.3	Max/min options	139
5.1.4	Quintessential option pricing formula	140
5.1.5	Symmetry results: Asian options	142
5.2	Affine jump diffusions/characteristic function methods	145
5.2.1	Lévy processes	145
5.2.2	Stochastic volatility	149
5.2.3	Pseudo-unification: affine jump diffusions	155
5.2.4	General results/contour integration	157
5.2.5	Specific examples	161
5.2.6	Application to change of measure	166
5.2.7	Spot and implied forward models	169
5.2.8	Fundamental drivers and exogeneity	174
5.2.9	Minimal martingale applications	178
5.3	Appendix	184
5.3.1	More Asian option results	184
5.3.2	Further change-of-measure applications	187
6	Econometric Concepts	191
6.1	Cointegration and mean reversion	191

6.1.1	Basic ideas	191
6.1.2	Granger causality	197
6.1.3	Vector Error Correction Model (VECM)	199
6.1.4	Connection to scaling laws	205
6.2	Stochastic filtering	207
6.2.1	Basic concepts	207
6.2.2	The Kalman filter and its extensions	209
6.2.3	Heston vs. generalized autoregressive conditional heteroskedasticity (GARCH)	220
6.3	Sampling distributions	225
6.3.1	The reality of small samples	225
6.3.2	Wishart distribution and more general sampling distributions	226
6.4	Resampling and robustness	231
6.4.1	Basic concepts	231
6.4.2	Information conditioning	232
6.4.3	Bootstrapping	235
6.5	Estimation in finite samples	237
6.5.1	Basic concepts	237
6.5.2	MLE and QMLE	242
6.5.3	GMM, EMM, and their offshoots	244
6.5.4	A study of estimators in small samples	247
6.5.5	Spectral methods	255
6.6	Appendix	258
6.6.1	Continuous vs. discrete time	258
6.6.2	Estimation issues for variance scaling laws	260
6.6.3	High-frequency scaling	268
7	Numerical Methods	272
7.1	Basics of spread option pricing	272
7.1.1	Measure changes	272
7.1.2	Approximations	275
7.2	Conditional expectation as a representation of value	279
7.3	Interpolation and basis function expansions	279
7.3.1	Pearson and related approaches	280
7.3.2	The grid model	285
7.3.3	Further applications of characteristic functions	300
7.4	Quadrature	304
7.4.1	Gaussian	305
7.4.2	High dimensions	313
7.5	Simulation	318
7.5.1	Monte Carlo	319
7.5.2	Variance reduction	323

7.5.3	Quasi-Monte Carlo	333
7.6	Stochastic control and dynamic programming	337
7.6.1	Hamilton-Jacobi-Bellman equation	338
7.6.2	Dual approaches	338
7.6.3	LSQ	339
7.6.4	Duality (again)	344
7.7	Complex variable techniques for characteristic function applications	346
7.7.1	Change of contour/change of measure	346
7.7.2	FFT and other transform methods	353
8	Dependency Modeling	359
8.1	Dependence and copulas	359
8.1.1	Concepts of dependence	359
8.1.2	Classification	365
8.1.3	Dependency: continuous vs. discontinuous processes	374
8.1.4	Consistency: static vs. dynamic	376
8.1.5	Wishart processes	381
8.2	Signal and noise in portfolio construction	383
8.2.1	Random matrices	384
8.2.2	Principal components and related concepts	389
	<i>Notes</i>	391
	<i>Bibliography</i>	437
	<i>Index</i>	451

List of Figures

1.1	Comparison of volatilities across asset classes	2
1.2	Spot electricity prices	2
1.3	Comparison of basis, leg, and backbone	7
2.1	AR(1) coefficient estimator, nearly non-stationary process	24
2.2	Distribution of t -statistic, AR(1) coefficient, nearly non-stationary process	25
2.3	Components of AR(1) variance estimator, nearly non-stationary process	26
2.4	Distribution of t -statistic, AR(1) variance, nearly non-stationary process	26
2.5	Illustration of non-IDD effects	38
2.6	Monthly (average) natural gas spot prices	39
2.7	Monthly (average) crude oil spot prices	40
2.8	Variance scaling law for spot Henry Hub	40
2.9	Variance scaling law for spot Brent	41
2.10	QV/replication volatility term structure, natural gas	41
2.11	QV/replication volatility term structure, crude oil	42
2.12	Front month futures prices, crude oil, daily resolution	43
2.13	Front month futures prices, natural gas, daily resolution	43
2.14	Brent scaling law, April 11–July 14 subsample	44
2.15	Henry Hub scaling law, April 11–July 14 subsample	44
2.16	Average Boston area temperatures by month	45
2.17	Variance scaling for Boston temperature residuals	45
2.18	Representative market heat rate (spot)	46
2.19	Variance scaling law for spot heat	47
3.1	Comparison of variance scaling laws for different processes	60
3.2	Expected value from different hedging strategies	63
3.3	Realized (pathwise) heat rate ATM QV	65
3.4	Comparison of volatility collected from different hedging strategies	66
3.5	Volatility collected under dynamic vs. static strategies	66
3.6	Comparison of volatility collected from static strategy vs. return volatility	67
3.7	Static vs. return analysis for simulated data	68
3.8	Typical shape of natural gas forward curve	73
3.9	Comparison of cash flows for different storage hedging strategies	74
3.10	Valuation and hedging with BS functional	97
3.11	Valuation and hedging with Heston functional	97
3.12	Portfolio variance comparison, EMM vs. non-EMM	98

3.13	Comparison of volatility projections	99
4.1	Implied daily curve	120
4.2	Daily and monthly values	121
4.3	Bounded tolling valuations	127
5.1	Contour for Fourier inversion	158
5.2	Volatility term structure for mixed stationary/non-stationary effects . . .	166
5.3	Volatility term structure for static vs. dynamic hedging strategies	174
5.4	Volatility modulation factor for mean-reverting stochastic mean	188
5.5	Forward volatility modulation factor for stochastic variance in a mean-reverting spot model	189
6.1	OLS estimator, “cointegrated” assets	193
6.2	OLS estimator, non-cointegrated assets	194
6.3	Standardized filtering distribution, full information case	220
6.4	Standardized filtering distribution, partial information case	220
6.5	Distribution of t -statistic, mean reversion rate	250
6.6	Distribution of t -statistic, mean reversion level	250
6.7	Distribution of t -statistic, volatility	251
6.8	Distribution of t -statistic, mean reversion rate	252
6.9	Distribution of t -statistic, mean reversion level	252
6.10	Distribution of t -statistic, volatility	253
7.1	Comparison of spread option extrinsic value as a function of strike	277
7.2	Comparison of spread option extrinsic value as a function of strike	278
7.3	Convergence rates, grid vs. binomial	296
7.4	Grid alignment	299
7.5	Convergence of Gauss-Laguerre quadrature for Heston	306
7.6	Convergence results for 2-dimensional normal CDF	311
7.7	Convergence of Gaussian quadrature	315
7.8	Convergence of Gaussian quadrature	315
7.9	Delta calculations	332
7.10	Comparison of greek calculations via simulation	332
7.11	Clustering of Sobol’ points	336
7.12	Sobol’ points with suitably chosen seed	336
7.13	Convergence of quasi- and pseudo-Monte Carlo	337
7.14	Integration contour for quadrature	352

List of Tables

3.1	Typical value drivers for selected energy deals	79
4.1	Daily and monthly values	120
4.2	Representative operational characteristics for tolling	126
4.3	Representative price and covariance data for tolling	126
7.1	Runtimes, grid vs. binomial	295
7.2	Comparison of quadrature techniques	318
7.3	Importance sampling for calculating $\Pr(z > 3)$ for z a standard normal	347
7.4	Quadrature methods for computing $\Pr(z > 3)$ for z a standard normal	349
7.5	Quadrature results for standard bivariate normal	351
7.6	Comparison of OTM probabilities for Heston variance	352

Preface

Energy markets (and commodity markets in general) present a number of challenges for quantitative modeling. High volatilities, small sample sizes, structural market changes, and operational complexity all make it very difficult to straightforwardly apply standard methods to the valuation and hedging of products that are commonly encountered in energy markets. It cannot be denied that there is an unfortunate tendency to apply, with little skeptical thought, methods widely used in financial (*e.g.*, bond or equity) markets to problems in the energy sector. Generally, there is insufficient appreciation for the trade-off between theoretical sophistication and practical performance. (This problem is compounded by the temptation to resort to, in the face of multiple drivers and physical constraints, computational machinations that give the illusion of information creation through ease of scenario generation *i.e.*, simulation.) The primary challenge of energy modeling is to correctly adapt what *is* correct about these familiar techniques while remaining fully cognizant of their limitations that become particularly acute in energy markets. The present volume is an attempt to perform this task, and consists of both general and specialized facets.

First, it is necessary to say what this book is *not*. We do not attempt to provide a detailed discussion of any energy markets or their commonly transacted products. There exist many other excellent books for this purpose, some of which we note in the text. For completeness and context, we provide a very high-level overview of such markets and products, at least as they appear in the United States for natural gas and electricity. However, we assume that the reader has sufficient experience in this industry to understand the basics of the prevailing market structures. (If you think a toll is just a fee you pay when you drive on the highway, this is probably not the right book for you.) Furthermore, this is not a book for people, regardless of existing technical ability, who are unfamiliar with the basics of financial mathematics, including stochastic calculus and option pricing. Again, to facilitate exposition such concepts will be introduced and summarized as needed. However, it is assumed that the reader has a reasonable grasp of such necessary tools that are commonly presented in, say, first-year computational finance courses. (If your first thought when someone says “Hull” is convex hull, then you probably have not done sufficient background work.)

So, who *is* this book for? In truth, it is aimed at a relatively diverse audience, and we have attempted to structure the book accordingly. The book is aimed at readers with a reasonably advanced technical background who have a good familiarity with

energy trading. Assuming this is not particularly helpful, let us elaborate. Quantitative analysts (“quants”) who work on energy-trading desks in support of trading, structuring, and origination and whose job requires modeling, pricing, and hedging natural gas and electricity structures should have interest. Such readers should have the necessary industry background as well as familiarity with mathematical concepts such as stochastic control. In addition, they will be reasonably expected to have analyzed actual data at some point. They presumably have little trepidation in rolling up their sleeves to work out problems or code up algorithms (indeed, they should be eager to do so). For them, this book will (hopefully) present useful approaches that they can use in their jobs, both for statistical work and model development. (As well, risk control analysts and quantitatively oriented traders who must understand, at least at a high level, valuation methodologies can also benefit, at least to a lesser extent.)

Another category of the target audience is students who wish not only to understand more advanced techniques than they are likely to have seen in their introductory coursework, but also to get an introduction to actual traded products and issues associated with their analysis. (More broadly, academics who have the necessary technical expertise but want to see applications in energy markets can also be included here.) These readers will understand such foundational concepts as stochastic calculus, (some) measure theory, and option pricing through replication, as well as knowing how to run a regression if asked. Such readers (at least at the student level) will benefit from seeing advanced material that is not normally collected in one volume (*e.g.*, affine jump diffusions, cointegration, Lévy copulas). They will also receive some context on how these methods should (and should not) be applied to examples actually encountered in the energy industry.

Note that these two broad categories are not necessarily mutually exclusive. There are of course practitioners at different levels of development, and some quants who know enough about tolling or storage, say, to operate or maintain models may want to gain some extra technical competency to understand these models (and their limitations) better. Similarly, experienced students may require little technical tutoring but need to become acquainted with approaches to actual structured products. There can definitely be overlap across classes of readership.

The structure of the book attempts to broadly satisfy these two groups. We divide the exposition into the standard blocks of theory and application; however, we reverse the usual order of presentation and begin with applications before going into more theoretical matters. While this may seem curious at first, there is a method to the madness (and in fact our dichotomy between practice and theory is rather soft, there is overlap throughout). As stated in the opening paragraph, we wish to retain what is correct about most quantitative modeling while avoiding those aspects that are especially ill-suited for energy (and commodity) applications. Broadly speaking, we present valuation of structured products as a replication/decomposition problem, in conjunction with robust estimation (that is, estimation that is not

overly sensitive to the particular sample). We essentially view valuation as a portfolio problem entailing representations in terms of statistical properties (such as variance) that are comparatively stable as opposed to those which are not (such as mean-reversion rates or jump probabilities). By discussing the core econometric and analytical issues first, we can more seamlessly proceed to an overview of valuation of some more popular structures in the industry.

In Part I the reader can thus get an understanding for how and why we choose our particular approaches, as well as see how the approaches manifests themselves. Then, in Part II the more theoretical issues can be investigated with the proper context in mind. (Of course, there is cross-referencing in the text so that the reader can consult certain ideas before returning to the main flow.) Although we advise against unthinkingly applying popular sophisticated methods for their own sake, it is unquestionably important to understand these techniques so as to better grasp why they can break down. Cointegration, for example, is an important and interesting idea, but its practical utility is limited (as are many econometric techniques) by the difficulty of separating signal from noise in small samples. Nonetheless, we show that cointegration has a relationship to variance scaling laws, which *can* be robustly implemented. We thus hope to draw the reader's attention to such connections, as well as provide the means for solving energy market problems.

The organization is as follows. We begin Part I with a (very) brief overview of energy markets (specifically in the United States) and the more common structured products therein. We then discuss the critical econometric issue of time scaling and how it relates to the conventional dichotomy stationarity/non-stationarity and variance accumulation. Next, we present valuation as a portfolio construction problem that is critically dependent on the prevailing market structure (via the availability of hedging instruments). We demonstrate that the gain from trying to represent valuation in terms of the actual *qualitative* properties of the underlying stochastic drivers is typically not enough to offset the costs. Finally we present some valuation examples of the aforementioned structured products.

Part II, as already noted, contains more theoretical material. In a sense, it fills in some of the details that are omitted in Part I. It can (hopefully) be read more profitably with that context already provided. However, large parts of it can also serve as a stand-alone exposition of certain topics (primarily the non-econometric sections). We begin this part with a discussion of (stochastic) process modeling, not for the purposes of valuation as such, but rather to provide a conceptual framework for being able to address the question of *which* qualitative features should be retained (and which features should be ignored) for the purposes of robust valuation. Next we continue with econometric issues, with an eye toward demonstrating that many standard techniques (such as filtering) can easily break down in practice and should be used with great caution (if at all). Then, numerical methods are discussed. The obvious rationale for this topic is that at some point in any problem, actual computations must be carried out, and we go over techniques particularly relevant

for energy problems (*e.g.*, stochastic control and high-dimensional quadrature). Finally, given the key role joint dependencies play in energy markets, we present some relevant ideas (copulas being chief among these).

We should point out that many of the ideas to be presented here are more generally applicable to commodity markets as such, and not simply the subset of energy markets that will be our focus. Ultimately, commodity markets are driven by final (physical) consumption, so many of the characteristics exhibited by energy prices that are crucial for proper valuation of energy structures will be shared by the broader class of commodities (namely, supply-demand constraints and geographical concentration, small samples/high volatilities, and most critically, volatility scaling). We will not provide any specific examples in, say, agriculture or metals, except to note when certain concepts are more widely valid. We will also employ the term “commodity” in a generic, plain language sense. (So, reader beware!)

Acknowledgments

I would like to thank Alexander Eydeland and an anonymous referee for their helpful comments on earlier drafts of this book. They have helped make this a much-improved product; any remaining flaws and errors are entirely mine. I would also like to thank Piotr Grzywacz, Mike Oddy, Vish Krishnamoorthy, Marcel Stäheli, and Wilson Huynh for many fruitful and spirited discussions on quantitative analysis. I must also express a special intellectual and personal debt to Krzysztof Wolyniec. This book arose from a number of projects we have collaborated on over the years, and could not have come into being without his input and insights. His influence on my thinking about quantitative modeling simply cannot be understated. I would also like to thank Swiss Re for their support, and SNL for their permission to use their historical data.

1] Synopsis of Selected Energy Markets and Structures

1.1 Challenges of modeling in energy markets

Although it is more than ten years old at the time of this writing, Eydeland and Wolyniec (2003, hereafter denoted by EW) remains unparalleled in its presentation of both practical and theoretical techniques for commodity modeling, as well as its coverage of the core structured products in energy markets.¹ We will defer much discussion of the specifics of these markets to EW, as our focus here is on modeling techniques. However, it will still be useful to highlight some central features of energy markets, to provide the proper context for the subsequent analysis.²

1.1.1 High volatilities/jumps

Energy markets are characterized by much higher volatilities than those seen in financial or equity markets. Figure 1.1 provides an illustration.

It is worth noting that the general pattern (of higher commodity volatility) has persisted even in the post-crisis era of collapsing volatilities across markets. In large part, this situation reflects the time scales associated with the (physical) supply and demand factors that drive the dynamics of price formation in energy markets. These factors require that certain operational balances be maintained over relatively small time horizons, and that the arrival of new information propagates relatively quickly. Demand is a reflection of overall economic growth as well as stable (so to speak³) drivers such as weather. Supply is impacted by the marginal cost of those factors used in the production of the commodity in question. A familiar example is the generation stack in power markets, where very hot or very cold weather can increase demand to sufficiently high levels that very inefficient (expensive) units must be brought online.⁴ See Figure 1.2. for a typical example.

The presence of high volatilities makes the problem of extracting useful information from available data much more challenging, as it becomes harder to distinguish signal from noise (in a sample of a given size). This situation is further exacerbated by the fact that, in comparison to other markets, we often do not have much data to analyze in the first place.

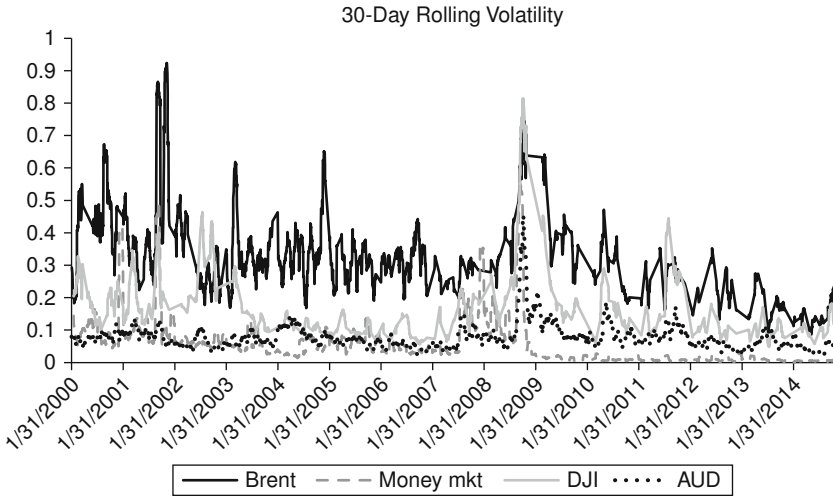


Figure 1.1 Comparison of volatilities across asset classes. Resp. Brent crude oil (spot), Federal funds rate, Dow Jones industrial average, and Australian dollar/US dollar exchange rate. Source: quandl.com.




Figure 1.2 Spot electricity prices. Source: New England ISO (www.iso-ne.com).

1.1.2 Small samples

The amount of data, both in terms of size and relevance, available for statistical and econometric analysis in energy markets is much smaller than that which exists in other markets. For example, some stock market and interest rate data go back to the early part of the 20th century. Useful energy data may only go back to the 1980s at best.⁵ This situation is due to a number of factors.

Commodity markets in general (and especially energy markets) have traditionally been heavily regulated (if not outright monopolized) entities (*e.g.*, utilities) and have only relatively recently become sufficiently open where useful price histories and time series can be collected.⁶ In addition (and related to prevailing and historical regulatory structures), energy markets are characterized by geographical particularities that are generally absent from financial or equity markets. A typical energy deal does not entail exposure to natural gas (say) as such, but rather exposure to natural gas in a specific physical location, *e.g.* the Rockies or the U.S. Northeast.⁷ Certain locations possess longer price series than others.

Finally, and perhaps most importantly, we must make a distinction between spot and futures/forward⁸ prices. Since spot commodities are not traded as such (physical possession must be taken), trading strategies (which, as we will see, form the backbone of valuation) must be done in terms of futures. The typical situation we face in energy markets is that for most locations of interest, there is either much less futures data than spot, or there is no futures data at all. The latter case is invariably associated with illiquid physical locations that do not trade on a forward basis. These include many natural gas basis locations or nodes in the electricity generation system. However, even for the liquidly traded locations (such as Henry Hub natural gas or PJM-W power), there is usually a good deal more spot data than futures data, especially for longer times-to-maturity.

1.1.3 Structural change

Along with the relatively recent opening up of energy markets (in comparison to say, equity markets), has come comparatively faster structural change in these markets. It is well beyond the scope of this book to cover these developments in any kind of detail. We will simply note some of the more prominent ones to illustrate the point:

- the construction of the Rockies Express (REX) natural gas pipeline, bringing Rockies gas into the Midwest and Eastern United States (2007–09)
- the so-called shale revolution in extracting both crude oil and natural gas (associated with North Dakota [Bakken] and Marcellus, respectively; 2010–present)
- the transition of western (CAISO) and Texas (ERCOT) power markets from bilateral/zonal markets to LMP/nodal markets (as prevail in the East; 2009–2010).

These developments have all had major impacts on price formation and dynamics and, as a result, on volatility. In addition, although not falling under the category of structural change as such, macro events such as the financial crisis of 2008 (leading to a collapse in commodity volatility and demand destruction) and regulatory/political factors such as Dodd-Frank (implemented after the Enron scandal in the early 2000s and affecting various kinds of market participants) have amounted

to kinds of regime shifts (so to speak) in their own right. The overall situation has had the effect of exacerbating the aforementioned data sparseness issues. The (relatively) small data that we have is often effectively truncated even more (if not rendered somewhat useless) by structural changes that preclude the past from providing any kind of guidance to the future.

1.1.4 Physical/operational constraints

Finally, we note that many (if not most) of the structures of interest in energy markets are heavily impacted by certain physical and operational constraints. Some of these are fairly simple, such as fuel losses associated with flowing natural gas from a production region to a consumer region, or into and out of storage. Others are far more complex, such as the operation of a power plant, with dispatch schedules that depend on fuel costs from (potentially) multiple fuel sources, response curves (heat rates) that are in general a function of the level of generation, and fixed (start-up) costs whose avoidance may require running the plant during unprofitable periods.^{9,10} Some involve the importance of time scales (a central theme of our subsequent discussion), which impact how we project risk factors of interest (such as how far industrial load can move against us over the time horizon in question).¹¹

In general, these constraints require optimization over a very complex set of operational states, while taking into account the equally complex (to say nothing of unknown!) stochastic dynamics of multiple drivers. A large part of the challenge of valuing such structures is determining how much operational flexibility must be accounted for. Put differently, which details can be ignored for purposes of valuation? This amounts to understanding the *incremental* contribution to value made by a particular operational facet. In other words, there is a balance to be struck between how much detail is captured, and how much value can be reasonably expected to be gained. It is better to have approximations that are robust given the data available, than to have precise models which depend on information we cannot realistically expect to extract.

1.2 Characteristic structured products

Here we will provide brief (but adequately detailed) descriptions of some of the more popular structured products encountered in energy markets. Again, EW should be consulted for greater details.

1.2.1 Tolling arrangements

Tolling deals are, in essence, associated with the spread between power prices and fuel prices. The embedded optionality in such deals is the ability to run the plant

(say, either starting up or shutting down) only when profitable. The very simplest form a tolling agreement takes is a so-called spark spread option, with payoff given by

$$(P_T - H \cdot G_T - K)^+ \quad (1.1)$$

with the obvious interpretation of P as a power price and G as a gas price (and of course $x^+ \equiv \max(x, 0)$). The parameters H and K can be thought of as corresponding to certain operational costs, specifically a heat rate and variable operation and maintenance (VOM), respectively¹² The parameter T represents an expiration or exercise time. (All of the deals we will consider have a critical time horizon component.)

Of course, tolling agreements usually possess far greater operational detail than reflected in (1.1). A power plant typically entails a volume-independent cost for starting up (that is, the cost is denominated in dollars, and not dollars per unit of generation),¹³ and possibly such a cost for shutting down. Such (fixed) costs have an important impact on operational decisions; it may be preferable to leave the plant on during uneconomic periods (*e.g.*, overnight) so as to avoid start-up costs during profitable periods (*e.g.*, weekdays during business hours). In general, the pattern of power prices differs by temporal block, *e.g.*, on-peak vs. off-peak. In fact, dispatch decisions can be made at an hourly resolution, a level at which no market instruments settle (a situation we will see also prevails for load following deals). There are other complications. Once up, a plant may be required to operate at some (minimum) level of generation. The rate at which fuel is converted to electricity will in general be dependent on generation level (as well as a host of other factors that are typically ignored). Some plants can also operate using multiple fuel types. There may also be limits on how many hours in a period the unit can run, or how many start-ups it can incur. Finally, the very real possibility that a unit may fail to start or fail to operate at full capacity (outages and derates, resp.) must be accounted for.

The operational complexity of a tolling agreement can be quite large, even when the contract is tailored for financial settlement. It remains the case, however, that the primary driver of value is the codependence of power and fuel and basic spread structures such as (1.1). The challenge we face in valuing tolling deals (or really any other deal with much physical optionality) is integrating this operational flexibility with available market instruments that, by their nature, do not align perfectly with this flexibility. We will see examples in later chapters, but our general theme will always be that it is better to find robust approximations that bound the value from below,¹⁴ than to try to perform a full optimization of the problem, which imposes enormous informational requirements that simply cannot be met in practice. Put differently, we ask: how much operational structure must we include in order to represent value in terms of both market information and entities (such as realized volatility or correlation) that can be robustly estimated? Part of our objective here is to answer this question.

1.2.2 Gas transport

The characteristic feature of natural gas logistics is flow from regions where gas is produced to regions where it is consumed. For example, in the United States this could entail flow from the Rockies to California or from the Gulf Coast to the Northeast. The associated optionality is the ability to turn off the flow when the spread between delivery and receipt points is negative. There are, in general, (variable) commodity charges (on both the receipt and delivery ends), as well as fuel losses along the pipe. The payoff function in this case can be written

$$\left(D_T - \frac{1}{1-f}R_T - K\right)^+ \quad (1.2)$$

where R and D denote receipt and delivery prices respectively, K is the (net) commodity charge, and f is the fuel loss (typically small, in the 1–3% range).¹⁵ Although transport is by far the simplest¹⁶ structure we will come across in this book, there are some subtleties worth pointing out.

In U.S. natural gas markets, most gas locations trade as an offset (either positive or negative) to a primary (backbone or hub) point (NYMEX Henry Hub). This offset is referred to as the basis. In other words, a leg (so to speak) price L can be written as $L = N + B$ where N is the hub price and B is the basis price. Thus, transacting (forward) basis locks in exposure relative to the hub; locking in total exposure requires transacting the hub, as well. Note that (1.2) can be written in terms of basis as

$$\left(B_T^D - \frac{1}{1-f}B_T^R - \frac{f}{1-f}N_T - K\right)^+ \quad (1.3)$$

Thus, if there are no fuel losses ($f = 0$), the transport option has *no* hub dependence. Hence, the transport spread can be locked in by trading in basis points *only*. Alternatively, (1.3) can be written as

$$\left(B_T^D - B_T^R - K - \frac{f}{1-f}R_T\right)^+ \approx (B_T^D - B_T^R - K)^+ - \frac{f}{1-f}R_T \cdot H(B_T^D - B_T^R - K) \quad (1.4)$$

We thus see that transport options are essentially options on a basis spread, and not a price spread as such. (Mathematically, we might say that a Gaussian model is more appropriate than a lognormal model.) Decomposing the payoff structure as in (1.4) we see that the optionality consists of both a regular option and a digital option, as well. We emphasize these points because they illustrate another basic theme here: market structure is critical for proper valuation of a product. Looking at leg prices can be misleading because in general (depending on the time horizon) the hub is far more volatile than basis. Variability in the leg often simply reflects variability in the hub. This is of course a manifestation of differences in liquidity,

which as we will see is a critical factor in valuation. For transport deals with no (or small) fuel costs, hedging (which is central to valuation through replication) will be conducted purely through basis, and care must be taken to not attribute value to hub variability.¹⁷ These points are illustrated in Figure 1.3.¹⁸ The implications here concern not simply modeling but (more importantly) the identification of the relevant exposure that arises from hedging and trading around such structures.

1.2.3 Gas storage

Another common gas-dependent structure is storage. Due to seasonal (weather-driven) demand patterns, it is economically feasible to buy gas in the summer (when it is relatively cheap), physically store it, and sell it in the winter (when it is relatively expensive). The embedded optionality of storage is thus a seasonal spread option:

$$\left((1 - f_{wdr})P_T^{wdr} - \frac{1}{1 - f_{inj}}P_T^{inj} - K \right)^+ \quad (1.5)$$

As with transport, there are typically fuel losses (on both injection and withdrawal), as well as (variable) commodity charges (on both ends, aggregated as K in (1.5)). However, unlike transport, there is no common backbone or hub involved in the spread in (1.5), and the underlying variability is between leg prices (for different temporal flows¹⁹).

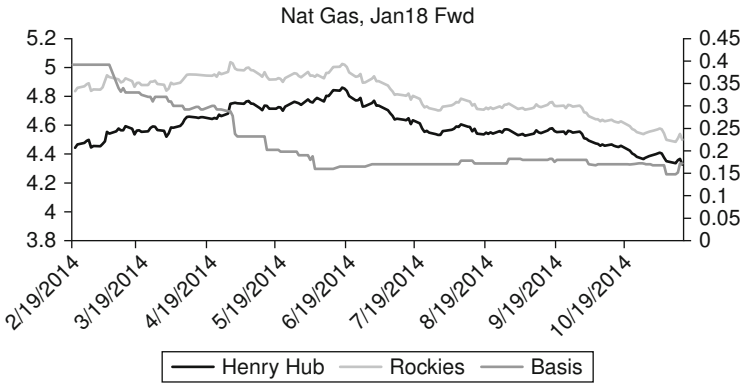


Figure 1.3 Comparison of basis, leg, and backbone. The (long-dated) Rockies (all-in) leg price clearly covaries with the benchmark Henry Hub price, but in fact Rockies (like most U.S. natural gas points) trades as an offset (basis) to Henry Hub. This basis is typically less liquid than the hub (esp. for longer times-to-maturity), hence the co-movement of Rockies with hub is due largely to the hub moving, and *not* because of the presence of a common driver (stochastic or otherwise). Source: Quandl.com.

One may think of the expression in (1.5) as generically representing the seasonal structure of storage. More abstractly, storage embodies a so-called stochastic control problem, where valuation amounts to (optimally) choosing how to flow gas in and out of the facility over time:

$$-\int_t^T q_s(f(q_s, Q_s)S_s + c(q_s, Q_s))d_s, \dot{Q} = q \quad (1.6)$$

where q denotes a flow rate (negative for withdrawals, positive for injections), Q is the inventory level, S is a spot price, and f and c are (action- and state-dependent) fuel and commodity costs, respectively. A natural question arises. The formulations of the payoffs in (1.5) and (1.6) appear to be very different; do they in fact represent very different approaches to valuation, or are they somehow related? As we will see in the course of our discussion, there is in fact a connection. The formulation in (1.5) can best be understood in terms of traded (monthly) contracts that can be used to lock in value through seasonal spreads, and in fact more generally through monthly optionality that can be captured as positions are rebalanced in light of changing price spreads (*e.g.*, a Dec–Jun spread may become more profitable than a Jan–Jul spread). In fact, once monthly volumes have been committed to, one is always free to conduct spot injections/withdrawals. We will see that the question of relating the two approaches (forward-based vs. spot-based) comes down to a question of market resolution (or more accurately the resolution of traded instruments). Put roughly, as the resolution of contracts becomes finer (*e.g.*, down to the level of specific days within a month), the closer the two paradigms will come.

As with tolling, there can be considerable operational constraints with storage that must be satisfied. The most basic form these constraints take are maximum injection and withdrawal rates. These are typically specified at the daily level, but they could apply over other periods as well, such as months. Other volumetric constraints are inventory requirements; for example, it may be required that a facility be completely full by the end of October (*i.e.*, you cannot wait until November to fill it up) or that it be at least 10% full by the end of February (*i.e.*, you cannot completely empty it before March). These kinds of constraints are actually not too hard to account for. A bit more challenging are so-called ratchets, which are volume-dependent flow rates (for injection and/or withdrawal). For example, an injection rate may be 10,000 MMBtu/day until the unit becomes half full, at which point the injection rate drops to 8,000 MMBtu/day. We will see that robust lower bound valuations can be obtained by crafting a linear programming problem in terms of spread options such as (1.5). The complications induced by ratchets effectively render the optimization problem nonlinear. As we stated with tolling, our objective will be to understand how much operational detail is necessary for robust valuation.

1.2.4 Load serving

The final structured product we will illustrate here differs from those we have just considered in that it does not entail explicit spread optionality. Load-serving deals (also known as full requirements deals) are, as the name suggests, agreements to serve the electricity demand (load) in a particular region for a particular period of time at some fixed price. The central feature here is volumetric risk: demand must be served at every hour of every day of the contract period, but power typically only trades in flat volumes for the on- and off-peak blocks of the constituent months. (Load does not trade at all.) Hedging with (flat) futures generally leaves one under-hedged during periods of higher demand (when prices are also generally higher) and over-hedged during periods of lower demand (when prices are also generally lower).

Of obvious interest is the cost-to-serve, which is simply price multiplied by load.²⁰ On an expected value basis, we have the following useful decomposition:

$$\begin{aligned} E_t L_{T'} P_{T'} &= E_t E_T (L_{T'} - E_T L_{T'} + E_T L_{T'}) (P_{T'} - E_T P_{T'} + E_T P_{T'}) \\ &= E_t [E_T (L_{T'} - E_T L_{T'}) (P_{T'} - E_T P_{T'}) + E_T L_{T'} \cdot E_T P_{T'}] \end{aligned} \quad (1.7)$$

Alternatively, we can write

$$\begin{aligned} E_t (L_{T'} - E_t L_{T'}) (P_{T'} - E_t P_{T'}) &= E_t E_T (L_{T'} - E_T L_{T'} + E_T L_{T'} - E_t L_{T'}) (P_{T'} - E_T P_{T'} + E_T P_{T'} - E_t P_{T'}) \\ &= E_t [E_T (L_{T'} - E_T L_{T'}) (P_{T'} - E_T P_{T'}) + (E_T L_{T'} - E_t L_{T'}) (E_T P_{T'} - E_t P_{T'})] \end{aligned} \quad (1.8)$$

In the expressions (1.7) and (1.8), t is the current time, T' is a representative time within the term (say, middle of a month), and T is a representative intermediate time (say, beginning of a month). These decompositions express the expected value of the cost-to-serve, conditioned on current information, in terms of expected values conditioned on intermediate information. For example, from (1.7), we see that the expected daily cost-to-serve (given current information) is the expected monthly cost-to-serve $E_t [E_T L_{T'} \cdot E_T P_{T'}]$ plus a cash covariance term $E_t [E_T (L_{T'} - E_T L_{T'}) (P_{T'} - E_T P_{T'})]$. (By cash we mean intra-month [say], conditional on information prior to the start of the monthly.) This decomposition is useful because we often have market-supplied information over these separate time horizons (*e.g.*, monthly vs. cash) that can be used for both hedging and information conditioning. (A standard approach is to separate a daily volatility into monthly and cash components.)

It is helpful to see the role of the covariance terms from a portfolio perspective. Recall that the deal consists of a fixed price P_X (payment received for serving the

load), and assume we put on a (flat) price hedge (with forward price P_F) at expected load (\bar{L}):

$$\Pi = -L_{T'} P_{T'} + L_{T'} P_X + \bar{L}(P_{T'} - P_F) = P_F \bar{L} \left(\frac{L_{T'}}{\bar{L}} - 1 \right) \left(\frac{P_{T'}}{P_F} - 1 \right) + (P_X - P_F) L_{T'} \quad (1.9)$$

Since changes in (expected) price and load and typically co-move, we see from (1.9) that the remaining risk entails both over- and under-hedging (as already noted). The larger point to be made, however, is that in many deals the (relative) covariation contribution $\left(\frac{L_{T'}}{\bar{L}} - 1 \right) \left(\frac{P_{T'}}{P_F} - 1 \right)$ covaries with realized price volatility.²¹ (This behavior is typically seen in industrial or commercial load deals [as opposed to residential]). Thus, an option/vega hedge (*i.e.*, an instrument that depends on realized volatility) can be included in the portfolio. As such, this relative covariation is not a population entity, but rather a *pathwise* entity. The fixed price P_X must then be chosen to not only finance the purchase of these option positions, but to account for residual risk, as well. Of course, this argument assumes that power volatility trades in the market in question; this is actually often not the case, as we will see shortly. However, in many situations one has load deals as part of a larger portfolio that includes tolling positions, as well, the latter of which have a “natural” vega component, so to speak. There is thus the possibility of exploiting intra-desk synergy between the two structures, and indeed, without complementary tolling positions load following by itself is not a particularly viable business (unless one has complementary physical assets such as generation [*e.g.*, as with utilities]).²² What we see here is a theme we will continue to develop throughout this book: the notion of valuation as a portfolio construction problem.

A final point we should raise here is the question of expected load. As already noted, load does not trade, so not only can load *not* be hedged, there are no forward markets whose prices can be used as any kind of projection of load. Thus, we must always perform some estimation of load. We will begin discussing econometric issues in Chapter 2 (and further in Chapter 6), but we wish to note here two points. First, the conditional decomposition between monthly and cash projections proves quite useful for approaching the problem econometrically. We often have a good understanding of load on a monthly basis, and can then form estimates conditional on these monthly levels (*e.g.*, cash variances,²³ *etc.*). Furthermore, we may be able to reckon certain intra-month (cash) properties of load robustly by conditioning on monthly levels. Second, load is an interesting example of how certain time scales (a central theme in this work) come into play. Some loads (residential) have a distinct seasonal structure, as they are driven primarily by weather-dependent demand. After such effects are accounted for, there is a certain residual structure whose informational content is a function of the time horizon in question. Currently, high or low demand relative to “normal” levels will generally affect our projections of future levels (again, relative to normal) inversely with time horizon.²⁴ On the other hand,

other load types (industrial or commercial) generally do not display sharp seasonal patterns, and are dominated by the responsiveness of customers to price and switching of providers (so-called migration). These loads (perhaps not surprisingly) have statistical properties more reminiscent of financial or economic time series such as GDP.²⁵ The informational content of such load observations accumulates directly with time horizon. We will give more precise meaning to these notions in Chapter 2.

1.3 Prelude to robust valuation

In this brief overview of energy markets and structures, we have already managed to introduce a number of important concepts that will receive fuller exposition in due course. Chief among these are the following facts:

- The perfect storm of small data sets, high volatility, structural change, and operational complexity make it imperative that modeling and analysis properly balance costs and benefits.
- Structured products do not appear *ab initio* but always exist within the context of a certain (energy) market framework that only permits particular portfolios to be formed around those structures.

These points are actually not unrelated. The fact that we have only specific market instruments available to us means that we can approach a given structured product from the point of view of replication or relative valuation, which of course means a specific kind of portfolio formation. Since different portfolios create different kinds of exposure, it behooves us to identify those portfolios whose resulting exposure entails risks we are most comfortable with.

More accurately, these risks are *residual* risks, *i.e.*, the risks remaining *after* some hedges have been put on.²⁶ Portfolio constructs or hedging strategies that require information that cannot be reliably obtained from the available data are not particularly useful, and must be strenuously avoided. We will have much more to say about valuation as portfolio formation in Chapter 3. Before that, we will first turn to the precursor of the valuation problem, namely the identification of entities whose estimation can be robustly performed given the data constraints we inevitably face in energy markets.