

2005

Risk Management in Long-Term Contracts

Victor P. Goldberg
Columbia Law School, vpg@law.columbia.edu

Follow this and additional works at: https://scholarship.law.columbia.edu/faculty_scholarship

Part of the [Business Organizations Law Commons](#), [Contracts Commons](#), and the [Law and Economics Commons](#)

Recommended Citation

Victor P. Goldberg, *Risk Management in Long-Term Contracts*, COLUMBIA LAW AND ECONOMICS WORKING PAPER No. 282 (2005).
Available at: https://scholarship.law.columbia.edu/faculty_scholarship/1382

This Working Paper is brought to you for free and open access by the Faculty Publications at Scholarship Archive. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarship Archive. For more information, please contact donnelly@law.columbia.edu.

Columbia Law School
The Center for Law and Economic Studies
435 West 116th St.
New York, NY 10027-7201

Working Paper No. 282

Risk Management in Long-Term Contracts

Victor P. Goldberg
September, 2005

*This paper can be downloaded without charge from the
Social Science Research Network electronic library at:*

http://ssrn.com/abstract_id=805184

*An index to the working papers in the Columbia Law
School Working Paper Series is located at:*

http://www.law.columbia.edu/center_program/law_economics

Risk Management in Long-Term Contracts

An Essay in Honor of Rudolf Richter's 80th Birthday

Abstract: Long-term contracts are designed to manage risk. After a brief discussion of why it is unhelpful to invoke risk aversion for analyzing serious commercial transactions between sophisticated entities, this paper focuses on adaptation to changed circumstances. In particular, it considers the options to abandon and the discretion to change quantity. It then analyzes a poorly designed contract between Alcoa and Essex showing how the parties misframed their problem and designed a long-term contract that was doomed to fail. (JEL: K3)

Professor Richter has played an important role in the revitalization of institutional economics. The conferences he ran for over a decade in Mettlach and Wallerfangen provided a forum for the growth of institutional economics and its spread beyond the U.S. borders. I attended a fair number of those conferences (including one on my honeymoon) and found them to be valuable both as an outlet for my work and as a stimulus. Indeed, two of my published papers grew out of discussions at the conference (Goldberg [1990, 1997]). My association with Professor Richter roughly coincided with my transition from an economics department to a law school. My new home has led me to focus more narrowly on matters of law than I had in the past. In recent years the focus has narrowed yet again to matters pertaining to contracts: the litigation side, the transactional side, and the regulatory side (that is, competition policy).

At one of Professor Richter's conferences in the late 1980's, Hal Varian presented a paper which purported to explain a particular practice (tax farming) by invoking risk aversion. I took issue with this and a somewhat heated discussion ensued. This eventually led to the publication of a brief paper in *JITE* explaining my "Aversion to Risk Aversion" (Goldberg [1990]). I had long been hostile to explanations of behavior that relied on particular assumptions about the risk preferences of parties, but this was the first opportunity I had to make the case in a systematic way. Perhaps I have been blinded by my strong views on risk aversion, but it is my impression that one of the distinguishing features of the New Institutional Economics has been a reluctance to base explanations of economic institutions on assumptions about risk preferences of individuals or organizations. Risk aversion should be, in my opinion, not the economist's first instinctive response; rather it should be the last resort.

Yes, of course, people are concerned about risk and contracts concern (among other things) the allocation of risks given that the future is uncertain. But the emphasis, especially when we are analyzing the decisions or behavior of sophisticated economic entities, should be not on risk *preferences*, but on risk *management*. If we ask, for example, why a public corporation, owned by shareholders who are (or could be) diversified, would buy actuarially unfair liability insurance, preference-based explanations require a lot of jerry-rigging.

It is more straight-forward to simply ignore risk preferences and ask why a risk-neutral corporation would purchase the package of services provided by a liability insurer. Those risk management services include inspection, litigation, and administration of compensation. If the insurance companies are more efficient at providing these services, other firms would pay them to do so, regardless of any attitudes toward risk. Would the insurance company charge a flat

fee? Probably not. The copayments and deductibles that are common in insurance contracts are plausible tools for dealing with the problem of assessing fault in a contract for services, particularly inspection. While the standard moral hazard argument would suggest that insurance inevitably increases the incidence and/or severity of accidents, the risk management perspective suggests that insurance—or, more precisely, the services provided by firms labeled as insurers—could reduce them.

I think that I have made some progress in my campaign against risk aversion, but the going is slow. One of the barriers is the lack of precision in the invocation of risk aversion. Even economists have more than one meaning. The standard economic definition has to do with risk preferences, the shape of the utility function, and the intuitive appeal has to do with the perception that most people seem to prefer certainty. But sometimes, especially when dealing with organizations, the economists will use risk aversion as a short-hand for any reason that might provide the proper curvature to whatever function they presume the entity happens to be maximizing. So, for example, if a big auto manufacturer shifts risks to its smaller suppliers, the formal economist would describe the big manufacturer as being relatively more risk averse without bothering to ask why. The institutional economist, at least this one, would bypass the uninformative risk-aversion label and instead ask why the small suppliers were the better bearers of the risk, and the focus would instead be on other factors like the relative costs and benefits of economizing on inventory holding.

While economists tend to blur the meaning, lawyers tend to obliterate it. Risk aversion often means nothing more than aversion to downside risk. That is, people as a rule don't like bad outcomes and they like really bad outcomes even less. Masochists aside, I am more than willing to concede that point. Imagine how difficult it is to carry on an anti-risk-aversion campaign when the target audience has adopted this bogus definition.

I should note that I am not making a claim about individual psychology. It might well be true that most people are risk averse most of the time. My point is that this is apt to be irrelevant for explaining contracting behavior, especially by large organizations like corporations. The art of theorizing lies in determining which simplifications of the world are useful for the class of questions one wants to focus on. Economists routinely make strong (and unrealistic) assumptions about the ability of individuals to process information and to perform complicated analytical tasks. As Milton Friedman reminded us over half a century ago, the expert billiards player can be modeled by assuming that he was applying the laws of physics even if he, in fact, was blissfully ignorant of those laws. Of course, the rules of physics apply equally to the terrible billiards player as well; it is just harder to view his game as if he were applying the laws of physics, rather than as the laws of physics determining what happens when he strikes the ball. My point is merely that to make sense of a complex world we need to make simplifying assumptions. For the part of the world I care about, a lot of mileage can be gained by assuming away attitudes toward risk.

That does not mean, of course, that we should assume away risk. The focus should be on mechanisms that cope with risk and uncertainty, mechanisms that can increase the expected value of a relationship. Let me illustrate by analyzing a very risky business decision, the structuring of a long-term supply contract. The entire discussion can be conducted without ever

assuming anything about the risk preferences of the firms or their executives. Assume that, for whatever reason, the parties have determined that a long-term deal would be better than either a series of spot transactions or common ownership. Regardless of the reason, the parties have to recognize when designing their relationship that the future is uncertain. What if demand shifts, technology changes, input prices change? Suppose the relationship proves disappointing? Could either party walk away or modify the terms? The value of the relationship could be enhanced if, as new information became available, the parties could readjust in pursuit of their common interest. The contract establishes a set of mechanisms for adaptation subject to the constraint that the parties are concerned not only with value creation, but with value division as well.

Suppose that after the parties enter into the agreement, the world unfolds and the seller realizes that continued performance would be unattractive. It would want to maintain the option to abandon, that is, to terminate. The right to cease performance if disappointing information is revealed during the course of the agreement is a valuable right. However, the seller's exercise of a right to abandon could impose a cost on the buyer who relied on the seller's continued performance. If the seller and buyer were a single entity, the abandonment decision would be a simple comparison of the costs and benefits. But they are not, by assumption, a single entity; the costs and benefits of abandonment have to be conveyed across organizational boundaries by the contract. The more the buyer intends to rely on the relationship, the higher the price the seller would have to pay to exercise its option. The option price could be set explicitly as a renewal option or in the form of liquidated damages; or the parties could just rely on the damage remedy for breach of contract. The option price need not be the same over all states of the world. Indeed, it is common to set the option price at zero under certain conditions. That is what *force majeure* clauses do.

The termination option is an extreme response to changed circumstances. A less extreme response would be to grant one party discretion to alter the arrangement as new information becomes available. For example, rather than establish a quantity schedule when the contract was formed, the contract could give one party discretion as to how much it would take in any given period. The contract would allocate the discretion to the party that valued it most. Suppose that in this instance that party was the buyer. The buyer's discretion would not be unconstrained, for if there were a tiny difference between the contract and market price the buyer could, in Judge Easterbrook's words, become the middleman to the world. (*Northern Indiana Public Service Company v. Colorado Westmoreland, Inc.* p. 636) Anyone foolish or careless enough to write such a contract deserves, I suppose, to go bankrupt, but typically, the buyer's discretion would be bounded by a physical constraint—all the coal to be used for running a particular power plant, for example. The physical constraint need not be the only limit on the buyer's discretion. The seller's reliance costs would again place limits on the buyer's discretion. So, for example, the buyer might have discretion in the form of a requirements contract. However, the contract could set a minimum or maximum (or both); if the buyer wanted to take a quantity outside that range, the contract could establish prices or a pricing mechanism that would force the buyer to take into account the consequences of its decision. And, as I mentioned in discussing the pricing of the option to abandon, the price can vary for different states of the world.

As an example of a contract where the price for exercising quantity discretion is low, consider an oil refiner contemplating adding a coker to its refinery, where there is little room for inventory accumulation of petroleum coke. I analyzed the structure of petroleum coke contracts written between 1946 and 1973 (Goldberg & Erickson [1987]). The coke is a byproduct of the production of higher valued products, like gasoline; as far as the refinery is concerned it is a waste product, but if the coke were to undergo a process called calcining, it would have value in the production of aluminum. If the coke accumulated at the refinery, it is possible that the refinery would have to shut down. The refiner does not want to have its output of gasoline and other high valued products dictated by its contract for disposal of the coke. Accordingly, the refiners typically entered into full output contracts with the owner of a calciner giving the refinery the discretion to produce as much, or as little, coke as it desired. The buyer in these contracts had the ability to hold a substantial amount of inventory, so the cost of allowing the seller such discretion was very low.

At the opposite extreme, consider the development of a coal mine and adjacent power plant, both with no practical alternative trading partners. The power plant would want the flexibility to vary its output in response to demand conditions, and would typically have a “requirements contract” which allowed the buyer to determine how much it would take, up to a contractually determined ceiling. However, the coal mine owner would want its substantial reliance to be taken into account. There are a number of mechanisms for conveying the seller’s reliance to the buyer. They are generally variations on a take-or-pay clause which establishes a nonlinear pricing scheme. The pure take-or-pay sets a fixed price even if nothing is taken, a zero marginal price up to the contract minimum, and a per unit price up to the contractual maximum quantity. The greater the mine’s reliance on this particular relationship, the more significant would be the limits on the power plant’s discretion.

To sum up, long term contracts often grant discretion over quantity to one of the parties. The discretion goes to the party that values it most; and that discretion is constrained by a pricing formula that conveys the counterparty’s expected costs when that discretion is exercised. This is not to say that long-term contracts would necessarily be for a variable quantity. Fixed quantity contracts are more likely if the buyer had easy access to resale markets or the seller had alternative suppliers. So, for example, a buyer might agree to take a fixed annual quantity knowing that it had the option of reselling part (or all) of its quota to others who were willing to pay a price higher than the buyer’s valuation.

The seller’s concerns about the buyer’s quantity discretion are exacerbated as the contract price becomes more out of line with current conditions. If the contract price were substantially below the current market price, the buyer’s incentive would be to maximize its requirements. If the requirements were for all the input to operate a particular plant, the buyer might take advantage by shifting production to this plant or by operating the plant beyond its planned capacity. The seller might counter by altering its performance by, for example, working to the rules or alleging the occurrence of force majeure events. The large contract-market differential would induce the seller to expend resources in attempting to renegotiate the terms of the agreement. The costs could arise directly from the effort to renegotiate or indirectly through strategic bargaining. That is, the loser might threaten to engage in acts which impose costs upon the other party but do not constitute a legal breach. If the parties had easy access to

market alternatives they would not be vulnerable to such strategic behavior. Other things equal, the more isolated from alternatives the contracting parties are, the more significant would be the potential losses. To the extent that the parties could anticipate these problems at the formation stage, they both bear the costs from this wasteful behavior. If the probability of wasteful behavior increases as the divergence between contract price and the opportunity costs of the parties widen, price-adjustment rules which narrow the gap become increasingly attractive.

I would suggest that this avoidance of the wasteful costs of renegotiation is a significant factor explaining the common use of price adjustment mechanisms in long-term contracts. This is not a matter of shifting exogenous risks to the less-risk-averse party. A high variance is costly, not because the parties are averse to risk, but because it leads to increased wasteful behavior. A properly designed price-adjustment mechanism could reduce that variance and thereby increase the expected value of a long-term deal.

Variance reduction can increase the expected value of the deal through a second channel. Suppose that the parties are contemplating entering into a multi-year contract. The contract establishes how the gains are to be divided between the parties. Each party could attempt to increase its share of the gains before signing the contract by improving its information on the future course of costs and prices. The more they each spend on this search, the smaller the pie. Other things equal, the larger the variance of the outcomes, the more resources would be devoted to this effort. The parties, therefore, have an incentive to incorporate into the initial agreement a device that would discourage this wasteful searching. One way to do that would be to replace a fixed price schedule with a price adjustment mechanism which could reduce the expected variance and the associated wasteful search.

I say “could” advisedly since it also “could not.” In some instances the parties would recognize that the adjustment mechanisms could exacerbate the problem. Perhaps there are no plausible indexes that would track market conditions. Perhaps including a trigger for price renegotiation would make one of the parties vulnerable to opportunistic reallocation of the gains from trade, or induce the parties to expend more resources in the renegotiation effort. These factors probably explain why aluminum companies operating calciners eschewed adjustment mechanisms in their contracts with coke producing oil refineries in the 1960's. See Goldberg & Erickson [1987, p. 391].

A second reason why the mechanisms “could not” reduce the expected variance is that they just bungled the deal. Economists like to explain observed behavior by presuming that the parties were behaving rationally. But we can also use economic reasoning to critique their behavior as well. In the remainder of this essay I want to analyze one such bungled deal, the long-term contract in which Alcoa sold molten aluminum to Essex, a fabricator of wire and cable. See Goldberg [2005, ch. 6.3].

Essex had been exploring a number of alternatives for assuring a supply of aluminum. These included building its own smelter, entering into a joint venture, or arranging a long-term supply agreement. Given the economies of scale in aluminum smelting as compared to fabrication, integration by ownership would have required that Essex become a seller of aluminum ingot. They settled instead on a long-term supply contract. I will simplify the terms

of the deal a bit. The contract was for 21 years for a fixed quantity of 75 million pounds per year. There were no restrictions on what Essex could do with the aluminum; that is, it was free to resell it rather than to use it. Essex's fabrication plant was to be located near Alcoa's Warrick, Indiana smelter so that it could take molten aluminum. This location enabled Essex to economize both on transportation costs and the costs of melting aluminum ingot. Alcoa had eight other smelters and the Essex contract accounted for less than one percent of Alcoa's annual production. Alcoa's reliance on this particular relationship was, therefore, trivial while Essex's was substantial.

The contract was designed to mimic what would have happened had Essex vertically integrated, while achieving the scale economies of Alcoa's smelter. Alcoa's chief executive testified: "Essex wanted a long term supply of aluminum . . . at a cost which would put them in, roughly, the same position as if they had gone into the primary aluminum business themselves . . . a situation where they could be paying our costs and a return on our investment, a reasonable return which they would have had to have calculated into their own economics if they had put their money into it" [Alcoa Brief, p. 6]. That is, to determine the contract price they took Alcoa's average capital cost of 40¢ per pound, calculated a reasonable rate of return on it, then added measures of Alcoa's labor and non-labor costs. Had this all been done well, Essex would have been in roughly the same position as if it had been given the ownership of a piece of Alcoa's smelter. The problems with this were twofold. It wasn't done well and it shouldn't have been done at all.

I'll begin with why it wasn't done well. The contract price was determined by a formula, starting with a base price of 15¢ per pound. There was a demand charge of 5¢ per pound and a component of production costs of 4¢ per pound; these were not indexed for the life of the contract. Labor costs (3¢ per pound) were indexed by Alcoa's average wage costs at the smelter, and non-labor production costs (also 3¢ per pound) were indexed by a component of the wholesale price index (WPI). Aluminum production is electricity-intensive and electricity accounted for much of Alcoa's non-labor costs. Alcoa's Warrick smelter relied on coal for its electricity. After the oil shock of 1973 the price of coal increased much faster than other components of the wholesale price index, so the index did a poor job of tracking non-labor production costs. Demand for aluminum increased as well. So by 1979, the market price of aluminum ingot was around 73 cents while Alcoa's costs were around 35 cents and the contract price was around 25 cents. Alcoa was not happy. It tried to get out of the deal or at least to revise it. After Essex refused, Alcoa brought suit asking for equitable reformation of the contract or, even better, termination. In a well-known opinion (*Alcoa v. Essex*), the district court, invoking the doctrines of mutual mistake, impracticability, and frustration of purpose, held for Alcoa and revised the terms of the contract. The crucial mistake, according to the court, was the decision by the two parties to use the wholesale price index to index non-labor costs.

That did turn out to be a bad decision, ex post, but it wasn't the worst. The more significant problem was that sixty percent of the initial contract price (the demand charge plus the fixed profit) was unadjusted for the life of the contract. A very simple example gives an indication of the type of problem this could cause. Suppose that the price level rises about 7% per year, doubling roughly every ten years; assume that the factors of production remain equally productive, that they continue to be used in the same proportions, and that the indexes for labor

and non-labor costs both worked perfectly. The indexed production costs would then rise from six cents per pound to 24 cents per pound in the twentieth year. However, since the remaining costs were unindexed, the final contract price would rise only to $(24 + 9 =)$ 33 cents. To keep the real price of aluminum constant the contract price would have had to increase to 60 cents. Even if the inflation rate had remained constant at its 1968 rate, by 1988 the contract price would have only risen to $(15 + 9) = 24$ cents while the nominal price of aluminum would have risen to 38 cents. Even that 14 cent differential was greater than the ten cent deviation that had moved the court to bail out Alcoa.

It should have been clear at the time of contracting, therefore, that the failure to index 60% of the price would inevitably yield a growing divergence between the contract and market price. Only if the real price of aluminum plummeted would this be avoided. If, as I argued above, price adjustment mechanisms are a way of reducing the gap between the market price and contract price, the contract was poorly designed to accomplish that. Even had the WPI faithfully tracked non-labor production costs, and even if the real price of aluminum had not changed, Alcoa would still have had strong incentives to attempt to renegotiate the price and to engage in behavior that would reduce the joint value of the deal.

Treating the construction costs and associated profits as sunk costs appears plausible, given the parties intent to treat Essex as the “as if” owner of a smelter. However, there is a problem with this line of reasoning. Suppose, to simplify, that Alcoa was building new capacity for its own use and had a target rate of return of 10%. That target rate assumes something about future rates of inflation. If the rate of inflation were to increase after the investment was made, the 10% return would be inadequate. However, if the inflation were a general phenomenon, the price at which Alcoa sold its aluminum would increase and the nominal rate of return would increase accordingly. So, for example, if the initial inflation rate were 3% and it was expected to remain at 3%, Alcoa’s nominal expected rate of return was 10% (by assumption) and its real expected rate of return would have been 7%. If the inflation rate doubled, if Alcoa continued to earn a nominal rate of 10%, the real return would have fallen to 4%. But if the inflation were general so that the real price of aluminum remained constant (that is, the nominal price of aluminum rose with the inflation rate), the real rate of return would have remained at 7% and the nominal rate would have risen to 13%. The point of this tedious exercise is that the failure to index the capital costs and rate of return meant that there was an implicit assumption that the rate of inflation would remain roughly unchanged for the duration of the agreement. To put this differently, in normal circumstances the capital costs would have been indexed implicitly by the demand side; the contract’s failure to index them meant that this component of costs would be inaccurately measured for the life of the contract.

So, had the parties better understood what they were doing, they could have produced a better price index. It still would have failed to track the market price, but at least it would not have been doomed. To cope with the likelihood that, over time, indexing to costs is likely to lose touch with market conditions, a common response would be to impose both a ceiling and a floor based on an external price. If either were triggered, there could be a number of possible responses. These include continuing to perform with the price remaining at the ceiling (or floor), termination, giving one party the option to terminate at a predetermined pricing formula, and having a third party set a new price. In fact, the Alcoa-Essex contract did include a ceiling as a

percentage of a published aluminum price. The ceiling is further evidence of their confusion. It was only eight percent above the initial contract price, but had inflation remained at the 1968 level and the real price of aluminum remained constant, the ceiling would have been 75% above the contract price by the end of the contract. It is hard to imagine a coherent economic reason for the ceiling to become increasingly less binding as time passed.

By framing their problem as one of treating Essex as if it owned the capacity to produce 75 million pounds of aluminum per year, therefore, the parties bungled the assignment of price risk. It also led them to err in assigning the quantity risk. The contract in effect made Essex the owner of capacity to produce 75 million pounds of aluminum per year. If it failed to use that quantity, Essex was still responsible for paying for it. It could either resell the unused portion or simply pay as if it had taken the aluminum. There was no good reason, however, for putting the quantity risk on Essex. Alcoa had other buyers for aluminum produced at Warrick. Its ingot plant at this location required more aluminum than could be produced at Warrick, so it was, on net, an importer at this location. Since sales to Essex under the contract amounted to less than 1% of Alcoa's annual sales, Alcoa had no reason to require significant protection of its reliance. While there was little reason to provide protection for Alcoa's reliance, that was not the case with Essex. The economics of its plant depended on the availability of molten aluminum. A more sensible contract would have given Essex the discretion to adapt to changing market conditions, probably in the form of a requirements contract for as much aluminum as necessary for fabrication at its nearby plant. Possibly, Alcoa might have wanted to include a minimum in order to limit Essex's ability to use the contract as an option—producing wire at this plant if the price here was better than at Essex's other plants and ceasing production here if the price elsewhere became more favorable. But the crucial point is that only one party had a significant reliance interest to protect, and the contract managed to get it backwards.

The Alcoa-Essex contract provides a nice example of how sophisticated parties can go wrong when they misstate their problem. Mimicking ownership might have seemed a plausible way to proceed, but it resulted in a badly designed deal, one that was likely to fail. The likelihood of failure was enhanced by the erroneous notion that the capital costs should not be indexed. Moreover, it led the parties to ignore the critical question: who should have the discretion to adapt to changed circumstances?

I have chosen to focus here on an instance in which the transactional lawyers asked the wrong question. I could instead have considered a number of cases in which the litigators failed to frame the question properly and the courts accepted their version. In particular, cases in which the courts imposed a "good faith" constraint on a requirements contract with no understanding of the rationale for that contract are common. A number of such cases are collected in Goldberg [2002]. The point of this exercise is twofold. First, the framework of the new institutional economics, in general, and transaction cost economics, in particular, can be extremely useful for analyzing behavior and influencing legal analysis. And, second, we can do just fine without invoking risk aversion.

Bibliography

Alcoa v. Essex 499 F.Supp. 53 (W.D.Pa. 1980)

Goldberg, Victor P. [1990]“Aversion to Risk Aversion in the New Institutional Economics,” *J. of Institutional and Theoretical Economics*, 146, 216-222.

_____ [1997] “The Net Profits Puzzle,” *Columbia Law Review*,97, 524-550.

_____ [2002]“Discretion in Long-Term Open Quantity Contracts: Reining in Good Faith,” *UCD Law Review* 35, 319-385.

_____ [2005] *Framing contract law: an economic perspective*, unpublished manuscript.

_____ and John Erickson. [1987] “Quantity and Price Adjustment in Long-Term Contracts: A Case Study of Petroleum Coke,” 30 *J. L. & Econ.* 30,369-398.

Northern Indiana Public Service Co. v. Colorado Westmoreland, Inc., 667 F.Supp. 613 (N.D. Ind. 1987).

Dr. Victor P. Goldberg
Columbia University
School of Law
435 W. 116th St.
New York, NY 10027
vpg@law.columbia.edu