

## Chapter 12: The Economics of Contract

I give you money, you give me an apple. No contract, no need for contract law.

I hire you to build a house on property I own. We agree on a price of a hundred thousand dollars. I give you a hundred thousand dollars and, in a world without enforceable contracts, never see you again.

The obvious solution is to make the payment due when the house is finished. You finish building my house and ask to be paid. I suggest that we renegotiate the terms. Until you are paid the house belongs to you, but it is on my land. If you do not want to accept my new and lower price you are free to tear it down again.

A better solution is to pay you continuously as you build the house, but that too has problems. When the house is three quarters built, you suggest renegotiating the price—upward. You have been paid for your work so far, and three quarters of a house is not of much use to me. I could pay someone else to finish it, but without the original contractor's detailed knowledge of just what has been done so far and what remains to be done, that may be a costly proposition.

Another solution, and a very common one, is reputation. You could cheat me but don't, because word will get around and nobody else will hire you to build houses. Reputation may be the most important method for enforcing agreements in our society, although not the one of most interest to lawyers.

You buy a sports jacket from a store that guarantees satisfaction—money back, no questions asked. When you unwrap it, your wife points out that you are a 42 short not a 40 long, and in any case purple is not really your color. You bring it back to the store and they give you your money back.

The store knows perfectly well that you are not going to sue them; the time and effort would cost you more than the jacket is worth. They give you your money back any way because they don't want a reputation for cheating their customers.

Department store refunds are an easy case, since the amounts at stake are small, the issues simple, and the parties are engaged in repeat dealings. Quite aside from what you may tell your friends, the store wants to keep you as a customer. Reputational enforcement is harder with larger amounts at stake and more complicated transactions. If your contractor agrees to be paid after the house is finished and you then insist on renegotiating the price down, you also take the precaution of thinking up some

plausible excuse, perhaps a list of complaints about exactly how the house was built. That way, third parties will be unsure who was trying to cheat whom. Even if they conclude that you are at fault, it may not matter much, since you are not planning to have another house built any time soon.

Even with large stakes, reputation sometimes provides an adequate mechanism for enforcing agreements. One such situation was explored in a classic article by Lisa Bernstein on the New York diamond industry. Buying and selling diamonds is a business where people routinely exchange large sums of money for envelopes containing lots of little stones without first inspecting, weighing, and testing each one.

At the time Bernstein studied the industry, it was dominated by orthodox Jews, forbidden by their religious beliefs from suing each other. What she observed was a trust intensive industry, an industry where the enforcement of agreements was of enormous importance, conducted almost entirely by people who could not use the legal system to enforce their agreements with each other. How was that possible?

The New York diamond markets had their own substitute for a court system: local rabbis with a reputation for sense and honesty. If one diamond merchant believed another had cheated him, he could take the dispute to the rabbi. If the rabbi ruled in his favor and the other party refused to pay up, he was out of the diamond business; nobody in the industry would trust him. The economic sanctions were reinforced by social sanctions, since the people who would conclude he was a crook would include most of his friends, relatives, and potential in-laws. Presumably the amount diamond merchants would be willing to risk on a single deal would depend in part on how long the other party had been involved in the industry and thus how much he would lose if he had to leave it.

As this account suggests, trusted private arbitrators provide an alternative to a court system, not by directly enforcing contracts but by generating the public information necessary to allow reputational mechanisms to enforce them. Third parties did not have to investigate a dispute in detail in order to learn who was at fault. They only had to find out which party the rabbi had ruled against. It also suggests a possible explanation for why particular trades are sometimes dominated by a single close knit ethnic group. If private enforcement of contracts via reputation and social pressure is less expensive and more reliable than enforcement via the court system, people whose social institutions make such private enforcement practical have a competitive advantage over those who must rely on the courts.

Reputation can do quite a lot to enforce contracts, but it cannot do everything. When I contract to have a house built, the odds are high that the contractor is not a member of the same close knit ethnic group that I am. If one of us violates our agreement the

other may not be in a position to prove the fact to interested third parties and thus impose reputational costs—either because the third parties do not know the victim exists or because the case is complicated enough so that it is hard for third parties to know who is at fault. And the violator, not being engaged in a trust intensive industry or repeat dealings, may regard the gain from breaking his agreement as worth any reputational cost it creates.

Even in this situation, private mechanisms may solve the problem, provided that we can agree on a third party we both trust. Each of us posts a bond with the third party. My bond forfeits if, in his judgement, I have broken the agreement, and similarly with yours. But while bonding permits private enforcement of agreements to cases that straight reputational enforcement does not reach, it still does not cover all situations.

My conclusion is that, although enforcement of contracts via courts and contract law is not the only way of solving these problems, and often not the best way, it may still have an important role to play. And whether or not it should play an important role, it in fact does, and is therefore worth studying. In any case, while the rest of this chapter is written in the context of contract law, most of the issues it discusses are issues of contract, not of law. They would arise, in only slightly different forms, for a New York rabbi or a commercial arbitrator.

## **Why Do We Need Contract Law?**

Given a system in which contracts are enforced by courts, why do we need contract law? Why not simply have the court read the contract and enforce the terms as read? That seems the obvious and just solution—and would save law students several semester hours.

There are at least three answers. The first is that courts may not wish to enforce contracts as written because they may believe that they know better than the parties what the terms should have been. A second is that even if you are willing to enforce the contract as written, you still have to decide whether a contract exists and what its terms are, questions about which the parties may sometimes disagree.

A third reason is that even if everybody agrees that a contract exists and should be enforced, and everyone agrees what it says, contracts never say enough. There is not enough fine print in the world to cover every possible contingency. Real-world

contracts cover some contingencies, typically ones the drafters think likely to occur, but leave gaps to be filled in by the court.

### **The Case for Freedom of Contract**

You and I are drawing up a contract for a joint business project. It occurs to me that a particular term is inefficient. Giving you an additional month to perform the first stage of the project would save you a hundred thousand dollars, cost me fifty thousand, and thus increase our net gain by fifty thousand dollars.

My first thought is that I should keep my mouth shut, since the change will leave me fifty thousand dollars worse off. But there is a better alternative. I propose to write the change into the contract along with a seventy-five thousand dollar increase in the amount you will pay me for my participation.

This simple example demonstrates a very general point, one that underlies the Coase theorem. As long as there exists a change in the terms of our deal that would produce net benefits, there is a way of making that change that benefits both of us. Hence we would expect rational bargainers to come up with contract terms that maximize the net gain. If our objective is economic efficiency, that is both an attractive outcome and an argument for enforcing contracts as written.

Courts frequently do enforce contracts as written, but not always. One reason they offer is "unequal bargaining power." The implicit metaphor is a world of conflict rather than cooperation, a war rather than a deal, in which only if each side is sufficiently powerful can it expect its concerns to be included in the final treaty.

That is a widespread and persuasive view of bargaining. To see what is wrong with it, consider a real-world case of unequal bargaining power. Suppose I am the only seller of this book—as, thanks to copyright law, I am. Further suppose that no alternative book is a close substitute for this one. That proposition too is true, as I hope to persuade you by the time you finish the last chapter. I am a monopolist. If you want a book like this, you must buy it on my terms or not at all.

Should I sell the book to you bound and with a cover, as books are usually sold, or simply put the unbound pages in a large envelope and sell that? Binding the book costs me money and provides benefits only for you. If a monopoly can impose its own terms without considering the desires of its customers, why should I bother?

The answer is that even a monopoly cannot force people to buy its products. As a rational monopolist, I am already charging you the highest price you are willing to pay—more precisely, I am already charging the price at which the cost of lost sales from any further increase would at least balance the gain from the higher price. Lowering the quality of the book will have the same effect as increasing the price. Some people (actually, lots of people) willing to pay my price for a bound book will be unwilling to pay the same price for an unbound one. As you may have noticed, this book came bound.

Suppose I think of some way of improving the book, such as a more detailed index or a better binding. The cost to me will be an extra two dollars per book, the increase in value to the customer will be three. It is in my interest to make the improvement and raise the price. The same would be true if instead of selling you a book I was renting you a car, and the improvement I was considering was not in the car but in the contract. If a change in terms is worth more to my customers than it costs me I should make it and adjust the price accordingly. We are back with our argument for efficient contract. Nothing in that argument depended on assuming that the seller was not a monopoly.

So far I have assumed that the change, in book or contract, has the same value to all customers. Suppose, however, that by improving the index and using higher quality paper I can increase the value of the book by five dollars to readers who really like it but only by one dollar to the more marginal customers. Again assume that the improvements cost me two dollars per book.

Even if the improvement is efficient—produces a net benefit—it is not in my interest to make it. What limits the price I can sell the book for is the value of the book to the marginal customers. The enthusiasts are already getting more than they are paying for.

Following out the logic of that example, we can see that improvements that selectively benefit marginal customers may be made even if not worth making, while improvements that selectively benefit non-marginal customers may not be made even if worth making. The result might be a level of quality higher or lower than would be efficient. That is a possible source of inefficiency and a legitimate argument against automatically accepting freedom of contract when one party is a monopoly. But it has very little to do with the conventional account of unequal bargaining power.

A second qualification to the conclusion that contracts will be drawn up with terms that maximize the net gain to the parties follows from my implicit assumption that the parties drawing up a contract are free to write into it any price that they can agree to. Suppose I am a landlord renting out an apartment whose market rent would be four hundred dollars a month. Unfortunately for me, the apartment is located in New York

City and covered by the city's rent control ordinance, which limits the rent to a maximum of three hundred dollars a month.

If the city controls the price but not the other terms of the contract, I respond by lowering the quality of the apartment in whatever ways save me the most money. As long as the value of the apartment is more than three hundred dollars, I will still be able to find tenants. By failing to keep the hot water system in good repair, or keep the building as warm as the tenants' would like it, or fix broken plumbing with reasonable speed, I am imposing a cost on the tenants which is greater than the savings to me; that is why, if I were free to set my own price, I would not be doing those things. But since I am not free to set my own price, lowering the quality of my product is, from my standpoint, the next best alternative.

It follows that one consequence of rent control, more generally of price control, is to subvert the mechanisms that normally keep the quality of products and the terms of contracts efficient, providing an argument for legal restrictions on both.

### **The Case against Freedom of Contract**

The argument I have just given implies that, with a few exceptions, contracts will be designed to maximize the net benefits of those who sign them. But efficiency is supposed to take account of costs and benefits to everybody. In many cases we can ignore effects on other people, for reasons sketched out in chapter 2; in a market society, contracting parties, like individual actors, are not ordinarily in a position to impose net costs on third parties without their permission. There are, however, some important exceptions.

When I take out a contract on you with a hit man, I am attempting to impose quite substantial costs on a third party. If my hit man misses, the court will not help me get my money back. Nor will it enforce a contract by which I agree not to testify against someone else's illegal acts. Nor will American courts enforce contracts in restraint of trade, such as an agreement by members of a cartel to maintain a common price. In all of these cases, the fact that a contract is signed is evidence that it benefits the parties signing it, but not that it produces net benefits.

Implicit in the argument for freedom of contract, as in most of the economic analysis of law, is the assumption that people who sign contracts are rational, that they know their own interest and act to get it. The argument breaks down where that assumption

is implausible, which is why courts will not enforce contracts made by children or lunatics.

A less convincing version of the same justification appears when a court refuses to enforce a contract on the grounds that one of the parties must have been incompetent, since otherwise he would not have agreed to those terms. Implicit in that claim is the court's belief that it knows the interest of the party better than he does, and thus is competent to judge his incompetence.

Consider an example, in current legal doctrine, where that belief may well be false:

I agree to deliver ten thousand customized widgets to you by January 10th, for which you will pay me. I further agree that if I fail to meet the deadline, I owe you damages of a hundred thousand dollars. I fail to deliver, you sue. Do you collect?

The answer depends on whether the court believes that a hundred thousand dollars is a reasonable estimate of the cost to you of my failure. If it does, the *liquidated damages* term in our contract is enforceable. But if the court believes that a hundred thousand dollars greatly overestimates the real cost, it may decide that the term is a "penalty clause" and as such unenforceable.

Two assumptions go into this policy: that the court can estimate the real cost of my failure well enough to identify a penalty clause and that penalty clauses are never efficient. The first may be wrong; the second surely is.

To see why, consider the analysis back in chapter 5 of the choice between property rules and liability rules. Under a liability rule, if I use your property without your permission I must pay damages reflecting the injury to you. Under a property rule, if I want to use your property I must first get your permission; if I don't, I suffer a punishment designed, not to measure the injury to you, but to make sure I won't use your property without your permission.

A penalty clause is a private version of a property rule. I have given you a property interest in my delivering by January 10th. If I want to deliver after that, I must buy your permission. If I neither fulfill my contract nor obtain your permission to void it, I suffer a penalty.

Just as in chapter 5, so here, there are good reasons why parties might sometimes prefer a damage rule to a liability rule. They may believe that the cost of moving resources to their highest value use through the market is cheaper than doing it through the court system—because the cost of the relevant market transactions is low, because the cost of using the court is high, or perhaps both.

Perhaps you and I expect to be engaged in many such transactions and so have an interest in maintaining our reputation as reasonable people to do business with. I am confident that if I have problems meeting the January 10th deadline, you will renegotiate on reasonable terms. Further suppose that neither of us has a high opinion of the competence of the court and both prefer spending our money on things other than lawyers. It may be perfectly reasonable for us to choose a penalty clause, in the expectation that it will rarely be invoked, over the alternative of giving the court a free hand to decide damages. Yet the same legal system that routinely enforces property rules created by judges and legislatures refuses to enforce property rules privately created by the people they will bind.

### **Making Sense of Duress**

A mugger catches you alone in a dark alley and offers you a choice: Give him a hundred dollars or he kills you. You reply that your life is well worth the price but unfortunately you are not carrying that much cash. He offers to take a check. When you get home, should you be free to stop payment? Should a contract made under duress be enforceable?

The argument in favor of enforceability is that if the contract is not enforceable, the mugger will refuse your check—or accept it and then make sure you can't stop payment by killing you and cashing the check before news of your death reaches the bank. Seen from that perspective, it looks as though even a contract made under duress produces benefits for both parties and so should be enforceable. You prefer paying a hundred dollars to being killed, he prefers receiving a hundred dollars to killing you. Where's the problem?

The problem is that making the contract enforceable makes offering people the choice between their money and their life a much more profitable business—most of us have more in our checking accounts than in our wallets. The gain from enforceability is a better chance, if you are mugged, to buy yourself free. It must be balanced against the higher probability of being mugged. It seems likely that the current legal rule, holding contracts made under duress unenforceable, is the efficient one.

But that may not be true under all circumstances. A peace treaty is a contract made under duress—yet most of us think that a world where nations can sign peace treaties and be bound by them is better than a world where the victor must annihilate the vanquished before he can be sure the war is over. Similarly, on a smaller scale, for the



transaction by which a prisoner of war gave his parole not to attempt to escape. Indeed, it used to be quite common for a prisoner to be released on parole and permitted to go home—having promised not to rejoin his army until he had been exchanged for a prisoner of equal rank from the other side. The parole system made war somewhat less costly for both sides and so presumably increased the amount of war somewhat, but it seems unlikely that the effect was very great—and it substantially decreased the cost born by captive and captor.

For an example of the same sort of tradeoff in a somewhat different context, consider a case that has recently been in the newspapers—the attempt by a Spanish judge to extradite Augusto Pinochet from England in order to try him in Spain for crimes he is accused of committing while dictator of Chile. Legal rules that immunize ex-dictators make it less expensive for them to commit crimes while in power. But legal rules that hold ex-dictators liable for such crimes make it more expensive for dictators to give up power. Pinochet is one of the rare examples of a dictator who voluntarily relinquished power to an elected government. If he ends up in a Spanish jail as a result, the next dictator may not make that mistake.

So far I have been discussing real duress: Your money or your life. There are other sorts of contracts that courts sometimes refuse to enforce on grounds of duress, but the duress is of a very different kind and the courts' decisions must be justified, if at all, with other arguments. We will consider two quite different examples, one that I like to think of as semi-real duress, the other as bogus duress.

### **Bargaining on a sinking ship**

Your ten million dollar ship has been caught in a storm, disabled, and is gradually going down. Fortunately, a tug comes by and offers to rescue it. Unfortunately, the tug captain, knowing the value of the ship, proposes to charge nine million dollars for his services. If you turn that offer down, he will be happy to take you and your crew to safety, leaving the ship to sink. You agree to his price, he tows the ship safe into harbor, and you refuse to pay, claiming that your agreement was obtained under duress. The admiralty court concludes that a reasonable price for the tow is one million dollars and rewrites your agreement accordingly.

Bargaining with the captain of the tug while the water is rising past your ankles certainly feels like duress to you. But this situation differs from real duress—your

money or your life—in one important way. The mugger got you into trouble; the tugboat gets you out. It was not his fault that your ship was sinking.

The argument offered earlier against enforcing contracts under duress here cuts in just the opposite direction. Making it easier for muggers to get your money increases the chance that someone will mug you, which is a bad thing. Increasing the amount tug boats can get for saving your ship increases the chance that, if you are sinking, a tug boat will be somewhere close, which is a good thing. If we are to justify the refusal to enforce such contracts, we need a different argument.

The first step to getting one is to ask what the efficient price is for the tugboat to charge—the price that maximizes the net gain to all concerned. Since the payment itself is merely a transfer, that means asking how the incentives created by the price will affect the actions of people who own tugboats and people whose ships might need a tug.

Consider the question first from the standpoint of the owner of the tugboat, deciding whether to spend an extra hundred thousand dollars to increase by one percentage point his chance of being at the right place in the right time to save a sinking ship. Perhaps he is deciding whether to cruise around in bad weather looking for ships in trouble, or whether to keep his radio receiver manned around the clock in the hope of picking up a distress call. When will he decide to spend the extra money and when, from the standpoint of efficiency, should he?

He will spend it if the price he can collect is at least ten million dollars, since in that case (assuming away any complications due to risk aversion) his hundred thousand dollar expenditure will produce an average return of at least a hundred thousand dollars. He should spend it if the value of the ship is at least ten million dollars, since in that case the social gain from his acts—an extra one percentage point chance of saving the ship—will be at least equal to the cost. If we want it to be in his interest to take those precautions that are worth taking and only those, we should allow him to collect the full value of the ship as the price for saving it. Any lower price means that, when he takes precautions, some of the benefit goes to the owner of the ship. That is a positive externality, and the result is a lower than optimal level of precautions.

Next consider the situation from the standpoint of the incentives of the ship owner. Suppose that, given whatever tug owners are doing, a sinking ship has a fifty percent chance of being rescued. The ship owner must decide what risks to take in running his ship—whether, for example, he should keep the ship in port during a storm or head out to sea, accepting a small risk that the ship may get into trouble. How will he make the decision and how, from the standpoint of efficiency, should he?

Sending the ship out in a storm has, we will suppose, a two percent chance of getting the ship into trouble. We further suppose, in the light of our analysis so far, that if the ship is rescued the charge for the rescue will be the full value of the ship. It follows that the cost of to the owner of sending the ship out is a one percent chance of the ship sinking plus another one percent chance of having to pay the tug ten million dollars to save it. The owner will send the ship to sea only if his benefit from doing so is at least two hundred thousand dollars.

This is the right calculation for him but the wrong calculation for an efficient outcome. Sending the ship to sea results in only one chance in a hundred of its sinking, since half the time it will be rescued. The rescue results in a ten million dollar loss to the ship owner but a matching gain to the tug owner. So it is efficient to send the ship to sea as long as the benefit is at least a hundred thousand dollars. To make it in the owner's interest to act that way, the price charged by the tug should be zero. As long as it is more than zero, the ship owner confers a positive externality on the tug owner every time he risks his ship, so he risks the ship less often than he should.

So far I have assumed away the actual cost to the tug in time, risk, and fuel of bringing the crippled ship back to port. If we redid the analysis including that, the conclusion would be that the price that gave the ship's owner the right incentive to keep his ship out of trouble was just equal to the actual cost of the rescue. Readers interested in a more precise statement of the argument and willing to put up with the necessary mathematics will find it on the web page.

[math icon]

We have just shown that there are two efficient prices. One, the full value of the ship, gives the tug owner the right incentive to be in a position to rescue sinking ships. The other, the cost of the rescue, gives the ship owner the right incentive to avoid getting his ship into a situation where it needs rescuing. Just where between those two the price should be set depends on how sensitive each party is to the relevant incentives. If there are very few acts the tug owner can take that would be worth taking at a price of ten million but not at a price of five, lowering the price to five million will cause only a small amount of inefficiency on the tow owner's side. Similarly, if there are very few cases where it would be worth putting out to sea if rescuers only charged their costs but not if they charged half the value of the ship, then a price of five million would cause only a small amount of inefficiency on the ship owner's side. By trading off such considerations one could, in principle, find the least bad price, the price that minimized the inefficiencies due to inadequate incentives on both sides.

The logic of the problem should be familiar from earlier chapters. The problem of two efficient prices, which I have gone into in some detail here, is simply Coase's problem

of joint causation applied to sinking ships instead of air pollution or airport noise. It will reappear in the next chapter in the context of auto accidents, which are also jointly caused. The reason it looks different here is that those examples involve dual causation of bad things, such as pollution or auto accidents, which we want to deter. This time we are concerned with dual causation of rescues, which we want to encourage. The solution is one we found back in chapter 7. Put the incentive where it will do the most good.

Our conclusion so far is that the efficient price is somewhere between the value of the ship and the cost of the rescue. A second conclusion is that there is no reason to expect bargaining to produce it. The bargaining occurs when the ship is sinking and the tug has already shown up, after all the relevant decisions, by ship owner and by tug owner, have already been taken. The only remaining decision is whether to rescue the ship or let it sink, and everyone already knows the right answer to that.

Everyone knows it, but we may not get it. When the bargaining occurs, it is in a setting of bilateral monopoly. Each side is trying to get as much as possible; while they argue, the ship is sinking. That suggests one good argument for the present legal rule, which permits an admiralty court to rewrite a contract that is too favorable to one side. It reduces the risk that the ship will sink while the two sides are haggling.

Alert readers may wonder how this analysis, in which the best possible result involves inefficiency on both sides, can be consistent with the Coase theorem. The answer is that in a fully Coasian world, a world where all transaction costs were zero, ship owners and tug owners would contract in advance to specify when tugs would be in places where ships might need rescuing, how willing ships would be to put to sea in storms, and the like. A sufficiently elaborate contract, in a world of zero transaction costs, would produce a result that was efficient on every margin. I have implicitly built transaction costs into my analysis by assuming that the contract is only negotiated after the ship gets into trouble. It would be interesting to investigate the actual market for salvage services in order to see where, between that simple picture and the fully efficient version, real arrangements end up.

### **Bogus Duress: Contracts of Adhesion**

A final category of duress accepted by some courts is of a sort familiar to all of us—form contracts, offered on a take it or leave it basis. When you rent a car from Avis you do not get to negotiate either the price or other terms of the contract. You are

presented with a series of options to take or refuse and, having done so, have the choice of either signing the contract or going to Hertz or Budget. Current legal jargon, borrowed from the French some decades back, refers to such contracts as contracts of adhesion.

This feels a little like duress. Certainly there is something one-sided about a transaction in which all of the terms are drawn up by one party. If your underlying view of contract formation is a conflict in which each party gets only the terms he bargains for, it is natural to suppose that the contract will end up as one-sided as the bargaining.

The reasons why this is wrong have already been discussed. Avis draws up the contract knowing that they still have to persuade you to sign it. They will, of course, write the most favorable contract to them that they think you will sign. But the way to do so is to draw up an efficient contract, thus maximizing the total gain, and then charge the highest price they think you will pay, transferring as much as possible of that gain to them.

We are left with the question of why Avis, Hertz, and many other firms, choose to make use of form contracts. There are two obvious reasons. One is that they want to minimize the costs of drawing up contracts. If you are making similar agreements with millions of customers a year, it is a lot cheaper to draft a single contract with options to cover likely variations among what your customers want than to redraft the contract for each transaction.

The second reason is that a form contract reduces the risk that their employees will cheat them. Suppose Avis left it up to each desk clerk to negotiate the rate of each rental. Five dollars to the desk clerk plus twenty dollars to Avis is a better deal for me than the thirty dollar price the clerk will hold out for if I don't bribe him. But it is a worse deal for Avis.

These arguments are relevant to the question of whether form contracts in general should be legally suspect; my conclusion is that they should not. There remains the possibility that in some cases form contracts may be suspect for special reasons—for instance, when the contract is so complicated that the customer does not know what he is signing.

That situation raises legitimate issues for a court trying to determine what the customer has agreed to; obviously terms printed in six point gray type and concealed in the contract's ornamental border do not qualify. On the other hand, one reason contracts written under our legal system are so elaborate may be the fear that courts will interpret any possible ambiguity as what the court thinks the contract should have

said instead of what the parties think it should have said; one way of protecting against that is to specify your contract in great detail. We will return to that problem, and what can be done about it, shortly. Similarly, one reason contracts may contain very one sided terms is that the parties are trying to limit the ability of courts to interpret the contract, while relying on nonlegal constraints, such as reputation, to prevent those terms from actually being enforced.

### **Contract, Contract, Who Has a Contract**

*"Otho would have been Bilbo's heir, but for the adoption of Frodo. He read the will carefully and snorted. It was, unfortunately, very clear and correct (according to the legal customs of hobbits, which demand among other things seven signatures of witnesses in red ink)."*

(J.R.R. Tolkien, *The Fellowship of the Ring*)

So far I have assumed the existence of a contract while considering various reasons why it might not be enforceable, including the possibility that agreement was not voluntary. A further set of issues arises in deciding whether or not there is a contract there to be enforced.

We agree to meet at a restaurant for dinner tomorrow. If you don't show up, am I entitled to sue you? Probably not. Agreements vary widely in how binding they are intended to be, and not all are properly the business of the legal system. We need ways of drawing lines, distinguishing between an unenforceable statement of intent and an enforceable contract, to enable people to communicate intentions without risking expensive entanglements with the legal system.

One solution is formality. For an agreement to count as a contract it must be signed in red ink, or sealed with purple sealing wax, or notarized, or deposited with the proper official, or all of the above. Such rules define a language, of actions rather than sounds, and so tell us what actions translate as "we intend this agreement to be enforceable."

A variety of other approaches are, in various legal systems, used to determine what is or is not a binding contract. One of the most important ones in our legal system is to

count as a contract only an agreement from which each party gains something—known as the *doctrine of consideration*. This raises an interesting problem:

My rich uncle, in an expansive mood at the family Christmas party, announces that a bright boy like me ought to go to college—and he is prepared to pay for it. Six months later, after being accepted by Harvard, I send him the bill: twenty-five thousand dollars for tuition, room and board. He sends it back, with a brief note explaining that he wasn't really serious. When I protest, he refers me to his lawyer, who patiently explains that since I didn't give him anything in exchange for his promise there was no consideration, hence there is no contract to enforce.

Fortunately for me, the situation is not that simple. I sue, invoking the doctrine of detrimental reliance. I did not give him anything in exchange for his promise, but I did spend quite a lot of time and effort, and some of my own money, getting into Harvard, as well as turning down several attractive job offers. These costs, all due to my relying on his promise, establish grounds for my claim that the promise was a contract and should be enforceable.

His promise imposed costs on me, so it makes some sense to hold him liable. It is less clear that contract law is the appropriate way of doing so, given the absence of most of the markers we usually use to tell what is or is not a contract. An alternative might be to count such rash promises as tortious. Another alternative is to impose no penalty at all, on the theory that before relying on his promise I should have gotten it in writing—notarized, sealed, and with the signatures of seven witnesses in red ink.

Interestingly enough, the requirement of consideration does not exist in property law. If I give you a piece of land as a gift, for no consideration, it is still yours; I don't get to change my mind next month and take it back. One possible explanation is that property law has other ways of determining whether a transfer has occurred; under the statute of frauds, all transfers of property must be in writing.

### **Contracts between fewer than two Parties:**

#### **Herein of Cats and Crash Victims**

Can there be a contract when one party does not know it exists? My cat gets out the door while my back is turned and I post a fifty dollar reward for her return. The next

day, a neighbor brings the cat back. The day after that, he happens to see one of my lost cat signs and calls up to demand the reward. Should he have a legal right to get it?

People who see a lost cat and do not know whether there is a reward will be more likely to try to find the owner if they know that, if a reward has been posted, they will be entitled to collect it. Their increased effort increases the chance that my cat will be found. This might in turn reduce efforts by finders who know about the reward, since the more likely someone else is to find the cat, the less the chance that your search will win the reward. But the combined effect must be an increased probability of finding my cat, since it is only if, on average, she is going to be found sooner that there is any reason for informed finders to reduce their effort.

[math link]

So far we have only considered the effect of the legal rule on the incentives of people who look for cats. What about its effect on the incentives of people who lose them? By making such contracts enforceable we make it more expensive to offer a reward; you may have to pay off even if the finder didn't know the reward has been offered until after he brought back the animal. On average, the result will be fewer offers of rewards. Whether the net effect is to increase or decrease the number of pets found and returned depends on the details of the situation: demand functions for lost cats, supply functions for potential searchers.

"'Contract' is not a preexisting entity of fixed dimensions; it is not a Platonic Form; it is the name given to a promise that courts will enforce. .... The answer to the question whether the return of the lost article is the acceptance of a contract offer should depend on whether, if so, more lost articles will be returned—a difficult question, as it happens, and one unrelated to logic." (Posner [1972] 1992: 251)

Posner's philosophical point is right but his economic conclusion is wrong. The return of lost articles, especially ones that purr, is a good. But time spent beating the woods for a lost cat in hope of reward is a cost. The relevant question is not whether more lost articles will be returned if such promises are enforced but whether the gain from the return of lost articles, net of the cost of searching for them, will be higher.

[case link]

For a final puzzle, consider the zero-sided contract, one that nobody has signed:



A physician comes upon an auto accident, stops, and treats an unconscious and badly bleeding victim. A week later, the victim receives a bill for services rendered. Must he pay it?

Under current United States law the answer is yes. To see why, consider again the analysis of the choice between property rules and liability rules. Under almost all other circumstances, the purchase of services is handled by a property rule—if I expect to provide you with services and be paid for doing so, I must first get your assent. I am not free to provide the services unasked and then bill you.

Why shouldn't billing for unasked benefits be treated as the flip side of suing for unasked injuries, a Pigouvian subsidy to compensate for a positive externality? The answer is that a property rule usually moves services to their highest valued use more cheaply and reliably than a liability rule. If the service is worth its price to me, you should be able to persuade me to buy it. That approach provides a cheaper and more reliable measure of the value of the services to me than providing services unasked and then suing for payment.

The accident victim, however, cannot contract for the service, because he is unconscious, even though the service is sufficiently valuable to make it almost certain that it is worth its cost to him. So in this case, unlike almost all others, courts enforce a negative liability rule, a Pigouvian bounty for the conferring of an unasked benefit.

### **Filling in the Blanks**

An opera star's employment contract covers many possible reasons why she might be unable to perform—but probably not the risk that she will be kidnapped by Martians or that a new religion will sweep the land and ban singing as a device of the devil. However carefully parties draw up contracts, there are always contingencies left out, either because nobody thought of them or because they were too unlikely to be worth including. When the unlikely becomes the certain, or when parties discover that they disagree about what the contract they agreed to means, the dispute may go to court. It is then up to the court to fill in the missing terms or resolve the ambiguity. How should it do so?

One plausible answer is that the court should try to figure out what the parties would have agreed to if they had covered the contingency or resolved the ambiguity. For

reasons discussed earlier in this chapter, that means that the court should try to find the efficient terms, the ones that maximize summed gains to the parties.

There are two arguments for doing this. The first is that efficiency is a good thing. The second is that such a policy reduces the cost of drawing up contracts. The parties can leave out unlikely contingencies, knowing that if they arise the court will try to fill in the terms they would have agreed to. One cost of the alternative policy of using omissions or ambiguities as an opportunity to implement the court's own agenda is that it gives the parties an incentive to write very long contracts.

If we accept this answer, the problem of interpreting contracts becomes the same as the problem of drafting them—figuring out what terms maximize the parties' summed benefit.

## **Risk Bearing**

I have agreed to deliver ten thousand customized widgets to you by January 10th. Early in the morning of January first, a drunk driver smashes his car through the wall of my warehouse, crushing half the widgets. One result is that I will have to replace them, at a cost of a hundred thousand dollars. A second is that I will deliver them a month late, costing you another hundred thousand dollars in lost sales. The drunk driver is dead and his estate bankrupt. Who pays for the losses? In chapter 6, we worked through in some detail the problem of efficiently allocating risk. It is now time to apply our results.

One basis for risk allocation is spreading losses. Suppose I am having a house built by a large firm that builds many houses each year. There is some risk that the house might burn down while it is being built. We could allocate that risk either to me, by specifying that I will have to pay for the additional construction costs in such a case, or to the builder. Since the builder is building many houses in different places, he can spread the risk; I cannot. From that standpoint at least, our contract should specify a fixed price, whether or not something goes wrong. If the contract does not specify who bears the risk, risk spreading provides an argument for assigning it to the builder.

A professional photographer spends six months taking photographs in the Himalayas for *National Geographic*, at a cost of a hundred thousand dollars. When he gets home, he gives his film to the local Walgreen's—which loses it. Do they owe him a hundred thousand dollars?

From the standpoint of risk spreading, the answer would be "yes." Walgreen's handles a large number of rolls of films each year, so can easily pool the risk. From the standpoint of moral hazard, incentives to keep the loss from happening, the answer is "no."

Walgreen's has no way of knowing that there is anything special about these rolls of film. The only way they can prevent the loss is by increased precautions on all rolls. They have the choice of an inefficiently low level of care for ten rolls of film or an inefficiently high level for ten million. The photographer does know that these films are especially valuable and can avoid the problem by taking them to a specialist film studio and making sure the proprietor realizes what they are. The efficient rule from the standpoint of moral hazard is to make the photographer liable, because he is the one in the best position to prevent the loss. That, called the rule of *Hadley v Baxendale* after an early case, is in fact the law.

In the case of the photographer, moral hazard and risk spreading cut in opposite directions; a firm drawing up a contract, or a court interpreting it, has to decide which effect is more important. In the case of a firm that builds houses the two arguments go in the same direction. The builder is not only in a better position than I am to spread the risk, he is also in a better position than I am to keep the house from burning down while he is building it.

A third factor relevant to allocating risk is adverse selection. Consider again the New Year's widget catastrophe. I did not know that that particular drunk driver was going to crash through that particular wall. But I probably do know a good deal more than you do about the chance that something—accident, strike, or bad planning—will prevent me from delivering the widgets to you on schedule. If our contract makes me liable for the resulting loss, you don't have to worry about that risk in deciding whom to buy your widgets from. You know what price I am charging, and you know that I am insuring you against the risk of non-delivery. If the contract specified that you had to bear the risk, you would need to know the reliability of alternative sellers as well as their price in order to decide which one is offering the best deal.

As this example suggests, moral hazard and adverse selection tend to cut in the same direction. As a general rule, the party with control over some part of the production process is in a better position both to prevent losses and to predict them. It follows that an efficient contract will usually assign the loss associated with something going wrong to the party with control over that particular something.

### **Getting Out of Contracts: Efficient Breach**

The year is 1929. I contract with you to clear land that I plan to develop in a few years; we agree on a price and a schedule. In drawing up the contract, we cover a number of problems that may arise—the land might flood, you might face a strike, I might get into a legal dispute over my title to the land. One of the contingencies we do not write into our contract is the great depression. Two years later, when you are about to start clearing, the real estate market has collapsed; the value of the cleared land will be less than the cost of clearing it. I tell you I don't want the land cleared; you respond that you have a contract to clear it and expect to be paid for doing so.

[case link]

As this example demonstrates, breaching a contract is not always a bad thing; there is little point to spending lots of money clearing land that nobody wants to build on. This raises an important question for contract law: Should breach be permitted, and if so what damages should the breaching party owe the victim of the breach?

The simplest legal rule is to impose no penalties at all. Some problems with that were discussed earlier in this chapter, in the context of *opportunistic breach*—the builder who takes his money and leaves or the buyer who renegotiates the price after the house is built. Others arise even if both parties honestly intended to fulfill the contract at the time they signed it.

We agree to a business deal which we expect will make each of us a hundred thousand dollars better off. After signing, something happens that raises my cost by a hundred and fifty thousand, converting a hundred thousand dollar gain into a fifty thousand dollar loss. Fulfilling the contract is still efficient, since your gain more than balances my loss. But it is a net loss for me so, in a world without enforceable contracts, I breach.

Or perhaps not. I inform you that I plan to breach the contract, and why. You respond by offering to pay me an additional sixty thousand dollars to fulfill the contract. Fulfilling now makes me ten thousand dollars better off and you forty thousand dollars better off, so there is no longer any reason to breach.

As this example suggests, the Coase theorem applies even in a world without enforceable contracts. As long as transaction costs are sufficiently low, contracts worth fulfilling will be fulfilled. The problem with unenforceable contracts is that they may lead to continuous costly renegotiation, with parties misrepresenting costs and gains in an attempt to shift the terms in their favor.

Next consider the opposite extreme, a legal regime in which breach is forbidden. A party to a contract is entitled to demand specific performance of the contract. Failure to perform results in severe penalties.

The situation is the same as before, except that this time my costs have risen by three hundred thousand dollars, converting a hundred thousand dollar gain into a two hundred thousand dollar loss. My loss is larger than your gain, so the efficient solution is to cancel the contract. But you are still better off if I fulfill the contract—and this time you have the legal right to make me do it. Fulfilling the contract is inefficient but profitable—for you—so you make me fulfill it.

Or perhaps not. I offer to pay you a hundred and fifty thousand dollars for permission to breach the contract (strictly speaking, it isn't breach once I have your permission, but if I pay attention to that distinction I will have to use "terminate," "rescind," and "breach" in different places to mean essentially the same thing, so I will continue to refer to all of them as "breach"). You accept. Breaching now makes each of us fifty thousand dollars better off than performing, so we breach. Generalizing the argument, we see that as long as transaction costs are low a rule of specific performance produces the efficient outcome. Contracts are breached if and only if performance produces a net loss.

Here again, problems arise because of transaction costs. When some unexpected problem makes it extremely costly for one party to perform its part of the contract, a rule of specific performance creates a bilateral monopoly bargaining problem. I am willing, if necessary, to pay you anything up to two hundred thousand dollars for permission to breach. You are willing, if necessary, to accept anything over a hundred thousand. The result may be a lot of costly bargaining, especially if costs and benefits are not as clear in the real world as in my example. Each side has an incentive to lie about its costs in order to get better terms—and to distrust the other side's claims. One result may be a breakdown of bargaining, resulting in the performance of a contract that it would be more efficient to breach.

Both unenforceability and specific performance are property rules in the sense discussed in chapter 5. With unenforceable contracts, each party has the right to breach, so performance happens only if both agree. Under specific performance, each party has the right to have the other perform, so breach happens only if both agree. In both cases, the mechanism for producing an efficient outcome, moving the rights to their highest valued use, is bargaining. And in both, such bargaining has the potential to produce large costs.

Unenforceable agreements are common—although we do not call them contracts. Specific performance is an uncommon rule except in contracts for the sale of real

property. Performance of such a contract is unlikely to be very inefficient; if the property I agreed to sell you is now worth more to me, and I cannot persuade you to cancel the contract, I can always sell the property to you and then buy it back.

The alternative to a property rule is a liability rule: Permit breach, but make the breaching party liable for damages. That makes sense if we believe that the costs of controlling breach by renegotiation are unacceptably high. But it also raises a new problem: How to set the damages so as to produce the efficient result.

The answer seems obvious. If I want to breach the contract, I must pay you enough to make you as well off as if I had fulfilled it—the Pigouvian solution to the problem of externalities. By making me liable, the law forces me to take account of all the costs due to my action. If my gain from breaching is more than your loss, I should breach and will. If not, I shouldn't and won't. The legal term for this rule is "*expectation damages*." If I breach, I must pay you enough to give you the result you expected from my performance.

The argument can be made clearer by going back to our earlier example. You expect to make a hundred thousand dollars from our contract. Under expectation damages, if I breach the contract I owe you a hundred thousand dollars. If fulfilling the contract costs me more than a hundred thousand I will breach and should. If it costs me less than a hundred thousand I won't breach and shouldn't.

This is the correct rule if our only concern is efficient breach, but (as usual) things are more complicated than that. A single legal rule affects incentives on a variety of different margins; we must take all of them into account in deciding on the right rule. One way of avoiding the costs associated with breaching a contract is to fulfill the contract. Another is not to sign it in the first place. How does the expectation rule affect the incentives relevant to that decision?

### **Getting Into Contracts Someone Might Want to Get Out Of: Efficient Signing**

Suppose that, if all goes well, signing the contract and fulfilling it makes each of us a hundred thousand dollars. If something goes wrong, signing the contract followed by my breaching it makes you a hundred thousand dollars better off and me a million dollars worse off. In deciding whether to sign the contract, I will take those facts into account.

The chance that something will go wrong, making breach efficient, is one in ten. Nine times out of ten I make a hundred thousand dollars, one time out of ten I lose a million. On average I am losing money so I don't sign the contract.

But, seen from our standpoint instead of my standpoint, I should sign it. Taking account of gains and losses to both parties, nine times out of ten we make two hundred thousand, one time out of ten we lose nine hundred. On average, we are making money. We need to revise the terms of the contract in a way that makes it in my interest as well as yours to sign it.

Doing so would be easy enough if both of us knew how likely each was to breach. Typically, I am better informed about events that could lead to my breaching than you are and you are better informed about events that could lead to your breaching than I am, which creates a problem. One solution is a rule under which each party only has to know the risk of his breaching in order to decide whether or not a contract is worth signing.

We get it by replacing expectation damages with *reliance damages*: If I breach I must make you as well off as if we had never signed the contract. I owe you compensation for any expenditures you made in reliance on my fulfilling the contract—hence the name of the damage rule—but I do not owe you for the profits you would have made if the contract had been fulfilled. Similarly, you owe me reliance damages if you breach the contract.

If I breach and then pay you reliance damages, you end up exactly as well off as if you had never signed the contract in the first place. So if I am going to breach it doesn't matter whether or not you sign. In order to decide whether to sign the contract, all you have to figure out is whether signing benefits you if I don't breach. In order to decide that you have to know how likely it is that you will breach, since that effects the risk that you will owe me damages, but not how likely it is that I will breach. Similarly, in order for me to decide whether to sign, I have to know how likely I am to breach, but not how likely you are.

Reliance is the correct rule if our concern is whether to sign contracts that might be breached, since it means that you do not have to know the risk that I will breach in order to know whether a contract is worth signing. It produces an efficient outcome on the sign/don't sign margin, assuming each of us knows the probability of events that will lead him to want to breach the contract, but an inefficient outcome on the breach/perform margin, since I will ignore your lost profits in deciding whether to breach. The expectation rule for calculating damages, on the other hand, produces efficient breach but inefficient signing. To put it differently, reliance damages force me to internalize the external cost produced by my signing a contract that I am, with

some probability known to me but not to you, going to breach. Expectation damages force me to internalize the external cost produced by my breaching a contract, given that it has already been signed.

### **Inefficient Reliance**

There is a third margin on which both rules produce an inefficient outcome. Imagine that you are deciding how to perform your part of the contract. One alternative is to spend nine hundred thousand dollars customizing your factory to produce the new product. Another is to use the factory as is, making production cost a million dollars higher than it would be in the customized factory.

If nothing is going to go wrong with the contract, you are better off customizing. Suppose, however, that there is one chance in five of some problem that will lead me to breach the contract and cancel my order—after you have customized but before you have produced. Four times out of five, customization saves a hundred thousand dollars. One time out of five it costs nine hundred thousand, since customizing a factory for a product and then not producing it is money wasted. On average, it is more efficient to use your existing factory.

But it is more profitable for you to customize. If I breach the contract I have to make you as well off as if I had not breached. Five times out of five, your decision to customize your factory gains you a hundred thousand dollars. One time out of five, it loses me nine hundred thousand.

Putting the argument more generally, under either rule potential victims of breach decide how much to rely as if the probability of being the victim of breach was zero—because if the other party breaches, the victim will be compensated for the money he wasted due to relying on the other's performance. But the efficient policy is to rely only if it is worth doing after allowing for the risk that something will go wrong. Both expectation damages and reliance damages result in an inefficiently high level of reliance.

There is a third alternative that solves that problem: Liquidated damages. Instead of having the court set liability on the basis of how much damage it thinks was done, the parties agree in advance on how much each will owe the other if he breaches the contract. Since the victim of breach gets the same amount whatever his reliance expenditures, liquidated damages result in an efficient level of reliance. If they are set



equal to what expectation damages would be, they also result in efficient breach; if they are set equal to what reliance damages would be, in efficient signing. It looks as though liquidated damages are unambiguously superior to both the other alternatives.

But liquidated damages must be set in advance, when the contract is signed. A court setting a damage payment after breach has happened has additional information about how much the real damage was. Thus whether liquidated damages are better or worse than the other alternatives depends in part on how serious a problem inefficient reliance is, in part on how well damages can be predicted in advance, and in part on how competent courts are at measuring damages after the fact.

### **Repeating Ourselves**

*"There is one story and one story only*

*That will prove worth your telling,"*

Robert Graves, "To Juan at the Winter Solstice"

If you think you see something familiar in the analysis of penalties for breach of contract, you are right. The basic arguments are the same as in the analysis of optimal insurance. Expectation damages are a solution to a problem of moral hazard, alias inefficient breach. Reliance damages are a solution to a problem of adverse selection, alias inefficient signing.

One of the problems with eliminating moral hazard by assigning risk to one party is that it increases moral hazard by the other, with the result that, in most cases, efficient risk allocation minimizes the inefficiency due to moral hazard but does not eliminate it. Making Coca-Cola liable for exploding bottles results in efficient precautions on the assembly line but inefficient explosions when someone is shaking a bottle of Coke to make a point at the fourth of July picnic. Similarly here. The loss due to my breach can be controlled by me by not breaching. It can be controlled by you by not relying too much on my performance. A legal rule that makes me liable for your losses due to my breach gives me the right incentive not to breach but gives you the wrong incentive to rely.

## The Boundaries of Fraud

*We fired our guns and the British kept a-coming*

*But there weren't near as many as there were a while ago;*

*We fired once more and the British started running*

*A-down the Mississippi to the Gulf of Mexico.*

### The Battle of New Orleans

The war of 1812 was fought in America but the peace treaty was signed in Europe, with the result that the most famous battle of the war was fought after the war was over. The same time delay that made Andrew Jackson a wholly unnecessary hero and, eventually, President also produced a famous case concerning the legal boundaries of fraud: *Laidlaw v Organ*.

Organ, a New Orleans merchant, got early word of the treaty of Ghent, which ended the war and, with it, the British blockade of New Orleans. He took advantage of the information by ordering a large quantity of tobacco from the Laidlaw firm. When the news of the treaty became public tobacco prices shot up. Laidlaw tried to renege on the contract. Organ sued, and won.

The economic argument against the verdict should be obvious: If the only reason Organ wants to buy the tobacco today is because it will be worth more tomorrow, the transaction produces no net benefit. Organ's profit comes entirely at Laidlaw's expense. He is engaged in rent seeking, and any resources he spends on that activity, such as money spent making sure he got the information a day early, are a net waste from the perspective of economic efficiency.

This example suggests the obvious counterargument. The particular use that Organ made of the information may have produced no net gain. But the production of information is itself a very valuable activity and the opportunity to profit by generating information thus provides a useful incentive. The sooner people in New

Orleans know the war is over, the sooner they can get back to the business of exporting tobacco.

Consider a commodity speculator who buys wheat today in the belief that it will be worth more next month. If he is correct, one result is that he makes money at the expense of whoever would otherwise have owned the wheat when its price went up. But another result is to drive up the price of wheat now, when he knows, and the rest of us don't know, that it is going to be scarce in the future. That gives other people an incentive to economize on their use of wheat, arrange for additional imports, and in a variety of other ways adjust in advance to future conditions of which they are as yet unaware.

The popular argument against speculation, that speculators, by keeping goods off the market, create shortages and price instability, is not merely wrong but backwards. A speculator who buys wheat today in order to sell it next month wants the price to be low this month and high next month, but the effect of his actions is precisely the opposite. He is driving the price of wheat up when it is plentiful, by buying it, and driving the price down when it is scarce, by selling it, thus smoothing out price fluctuations. The net effect is not to create shortages and famines but to prevent them, by providing ordinary people an incentive to adjust their present behavior to future conditions that can be foreseen by a sufficiently well qualified expert.

In order for this to work, the speculator must be able to profit by his superior knowledge. If the people who sold wheat to the speculator today were allowed to cancel the transaction after the price went up, there would be little point to being a speculator. If the Louisiana Supreme Court had ruled in favor of Laidlaw rather than Organ, and consistently followed the same rule in other cases, the result would have been less speculation and more unstable prices for agricultural commodities.

But while successful speculation is both productive and profitable, its productivity is not measured by how profitable it is. If Organ buys one day and sells the next, taking advantage of his early knowledge of the treaty, and if it happens that no decisions relevant to producing or using tobacco are made in that interval, he still makes money from the transaction. Thus speculation is an odd case of an economic activity that produces a private benefit that is matched by an external cost—the loss to the party that would have been holding the commodity when its price went up if the speculator had not bought it—along with an external benefit. The net effect is to provide speculators an incentive to produce valuable information and make that information public through the effect of their market transactions on prices, but that incentive is unrelated to just how valuable the information really is.

This analysis suggests a possible defense both for the decision in *Laidlaw v Organ* and for the willingness of courts in some other contexts to invalidate contracts on grounds of fraudulent concealment of information. We want a legal system in which people who acquire information can sometimes profit by doing so. But we do not want a system in which people who happen to have information highly relevant to the value of what they are selling—for example, the fact that the cow they are selling is afflicted with a serious and not easily observed disease, or the house they are selling has the reputation of being haunted—have an incentive to withhold it, thus producing transactions that may well produce a net loss rather than a net gain.

The haunted house is a real case. The plaintiff was attempting to cancel his purchase after discovering that the house he was buying was widely reputed to be infested by poltergeists. The judge ruled in his favor, at least in part because it was clear that the house's reputation was due to energetic publicity work by the seller. Having told everyone she could reach, including the subscribers of *Reader's Digest*, that her house was haunted, she was not entitled to withhold that information from a prospective purchaser.

"Whether the source of the spectral apparitions seen by defendant seller are parapsychic or psychogenic, having reported their presence in both a national publication (*Reader's Digest*) and the local press ... defendant is estopped to deny their existence and, as a matter of law, the house is haunted. ...

Finally, if the language of the contract is to be construed as broadly as defendant urges to encompass the presence of poltergeists in the house, it cannot be said that she has delivered the premises "vacant" in accordance with her obligation under the provision of the contract rider ... ."

Rubin, J. in *Stambovsky v. Ackely*