Introductory Remarks on Grossman and Hart, *JPE*, 1986

by

John Moore

You may be wondering: Who would have the temerity to make remarks about a masterpiece from the annals of economic thought? (Evidently, this piece started life as a speech, opening the conference to celebrate the twenty-fifth anniversary of the publication of Sandy and Oliver's article, and it hasn't managed to shake off its rhetorical roots.)

Let me explain. Back in 1986, the editors of the *Journal of Political Economy* in their wisdom published "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration" by Sanford Grossman and Oliver Hart. The same august journal subsequently published, in 1990, another article by Oliver Hart, "Property Rights and the Nature of the Firm", this time co-authored with John Moore. Taken together, the two articles make up what is known as the GHM Property Rights Theory. I am M. In recent Bond movies I've been played by both Judi Dench and Ralph Fiennes.

How did I get on to the GHM ticket? When I was a graduate student at the LSE, quite out of the blue Oliver phoned me up and asked if I would like to work with him. Me! I wasn't even his student. He wasn't even at the LSE. It was my Kate Middleton moment. Picked from a bevy of beauties by the Handsome Prince himself. Well, the parallels aren't exact. But I would like to take this opportunity to put on record that the moment Oliver chose to work with me was the turning point of my intellectual life, and I thank him.

I first met Sandy when he was visiting Oliver at the LSE. His reputation went before him: Completed a Chicago PhD at the age of 14¾, rapidly followed by tenure as an Ivy League Full Professor. When I knocked on Oliver's door, I couldn't believe it when Sandy sprang to his feet and warmly shook me by the hand saying what a pleasure it was for him to meet me. I think he mistook me for Jim Mirrlees.

The GHM moniker is so flattering (to me) that it's incumbent on me to propose a small modification. At the twenty-fifth anniversary celebrations, our drinks were laced with a drug having this molecular formula:

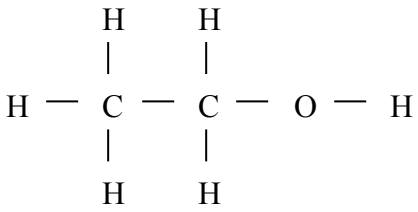

Ethyl alcohol: C_2H_6O. (Or, more properly, ethanol: CH_3CH_2OH.) No-one would ever write it as CHO, because that wouldn't give any sense of the respective contributions of the three elements that make up the molecule. And I think that by the same logic, no-one should write GHM. So I'd like to propose a citations-based way of writing the various multi-letter theories we have in our profession. I'm going to use a Google citation based weighting:

$$G_{5600}H_{9100}M_{3500}$$

That is one heck of a molecule. I should explain that the Grossman-Hart paper had just under 5600 Google citations at the time of the celebrations. And the Hart-Moore paper I'm afraid had only 3500. (The passage of time has been kind. At the time of writing (2014), these figures have risen to 8850 and 5114 respectively.)

If you divide by 700, you get something a little more manageable. So I propose that our theory is called the

$G_8H_{13}M_5$ Property Rights Theory

– to reflect the respective contributions of Grossman, Hart and Moore. Thank our good fortune it isn't known as "Grossman et al's Property Rights Theory", because I'd be Al and Oliver would be E.T.

On the subject of citations, which as you know is the only thing we are judged by, my second contribution here is to introduce what might be called the Average Moore Score (AMS) of any group of economists {i = 1, …, I}. The idea is to average over their highest number of Google citations:

$$\underset{\{\text{over i}\}}{\text{average}} \; \underset{\{\text{articles by i}\}}{\text{maximum}} \; [\text{number of Google citations for an article by i}]$$

To make it plain, take someone like me, and ask what is my best article in terms of Google citations – no surprize it's the 1990 paper with Oliver, which, at the time of the twenty-fifth anniversary celebrations, had 3500. Hence my contribution is 3500. Repeat the question for everyone in the group. Averaging then arrives at the group's AMS.

I did that for all the thirty seven Nobel Laureates in Economics since 1990, and, at the time of the celebrations, their AMS was 5533. Luckily, Grossman-Hart (1986) had 5576 Google citations, a bigger number. Just. Whew.

But note. That illustrious group's AMS, 5533, was dominated by the likes of Ronald Coase, and people who are cited by psychologists or statisticians. So, on reflection, I don't think averaging is necessarily the best way to calculate a Score. It may be better to look at the Median Moore Score (MMS):

$$\underset{\{\text{over i}\}}{\text{median}} \; \underset{\{\text{articles by i}\}}{\text{maximum}} \; [\text{number of Google citations for an article by i}]$$

The MMS for the same group of thirty seven Laureates was 2767. In that context, Grossman and Hart's 5576 looks good.

The *Journal of Political Economy* used to publish a Hall of Fame – sadly, no longer – which made up a formidable list of papers. For example: Hotelling (1931) on exhaustible resources; Tiebout (1956) on local expenditures; Conrad and Meyer (1958) on slavery; Stigler (1961) on the economics of information; Becker (1968) on crime and punishment; Barro (1974) on Ricardian equivalence; Kydland and Prescott (1977) on rules over discretion; Hall (1978) on the life cycle; Jovanovic (1979) on job matching; Fama (1980) on agency; Diamond and Dybvig (1983) on bank runs; Shleifer and Vishny (1986) on large shareholders; Romer (1986, 1990) on endogenous growth; Krugman (1991) on economic geography. These, and many other *JPE* papers, comprise an astonishing group. And, yet, according to the last *JPE*'s Hall of Fame ranking, which paper beat the lot? Grossman and Hart (1986)! I don't know how the ranking was arrived at, but the only *JPE* article that outranked Grossman-Hart (1986) – rhetorically, I could have done without this – is the Black and Scholes (1973) paper on options pricing (presumably the financial community added considerable heft).

Finally on the matter of statistics, we should always take into account Oliver's 1995 book. (I should say that the Moore Score, average or median, relates only to articles, not books.) Oh, that book. The Good Book. After W.Shakespeare there came O.Hart. *Firms, Contracts and Financial Structure*, a book for the masses: it just flies off the shelves of British bookstores. One can but feel for A.Smith and his *Wealth of Nations*.

Frank Hahn, on one of his royal visits down from Cambridge (England) to give a seminar to the peasants at the LSE, famously pronounced – and I remember this – "Grossman and Hart are now together in Cambridge working on the theory of the firm, but are very depressed about how hard it is." Since, according to Hahn, "Grossman and Hart have IQs unbounded above", it struck us graduate students that the theory of the firm must be very hard indeed.

I understand that the breakthrough in Sandy and Oliver's research came when Sandy acted as an expert witness in an anti-trust case. Naturally, Sandy was working pro bono for a little guy (AT&T) defending itself against an oppressive state. Sandy and Oliver found themselves bumping up against the most basic question of all: What is a firm? Indeed, that is the first sentence of their 1986 article. Four years later Oliver and I began our paper with the same opening line. We obviously must have felt that Sandy and he hadn't fully figured out the answer.

I would like to tell you about the answer they gave. This is Grossman and Hart (1986) viewed through the lens of Hart-Moore (1990) and Oliver's 1995 book. See Figure 1.

FIGURE 1 near here

There is pair of agents, managers M_1 and M_2 – say me and you respectively – and a pair of assets a_1 and a_2. There are three dates. At the earliest date 0, ex ante, the parties contract over asset ownership A_1 and A_2. At an interim date ½, a non-cooperative game is played of investments x_1 and x_2, which are private and non-contractible. At date 1, ex post, we may or may not trade (in equilibrium we will). If we trade, then our combined utility is $u_1(x_1)$ plus $u_2(x_2)$. Note that these utilities are functions of our respective date ½ investments. If we can't agree terms of trade (off the equilibrium path), then we go our separate ways, and I get $v_1(x_1|A_1)$. Notice that v_1 is a function both of my investment x_1 and, crucially, of the set of assets A_1 that I was given control over back at date 0 when we wrote our contract. And, in the event of breakdown in ex post bargaining, you get $v_2(x_2|A_2)$ – where A_2 is the set of assets over which you were given control.

I've already committed a serious crime by saying that the only thing we can contract over at date 0 is who gets to own – to have control over – which assets. The crime, of course, is not to justify why we don't write a contract at date 0 to govern the terms of trade at date 1. That is the way classical mechanism design would think about it. Even though I'm a criminal, let me plough on.

Since there are just the two of us, and just two assets, there is only a small set of simple ownership structures. Three candidates spring to mind. The first is *type 1 integration* where I own both assets and you don't own either. Next, *non-integration* where I own my asset, a_1, and you own yours, a_2. Finally, *type 2 integration* where you own both assets and I own nothing. (I believe Grossman and Hart (1986) was the first time anyone had formally made the point that there could be a difference between your acquiring my assets – type 2 integration – and my buying yours – type 1.)

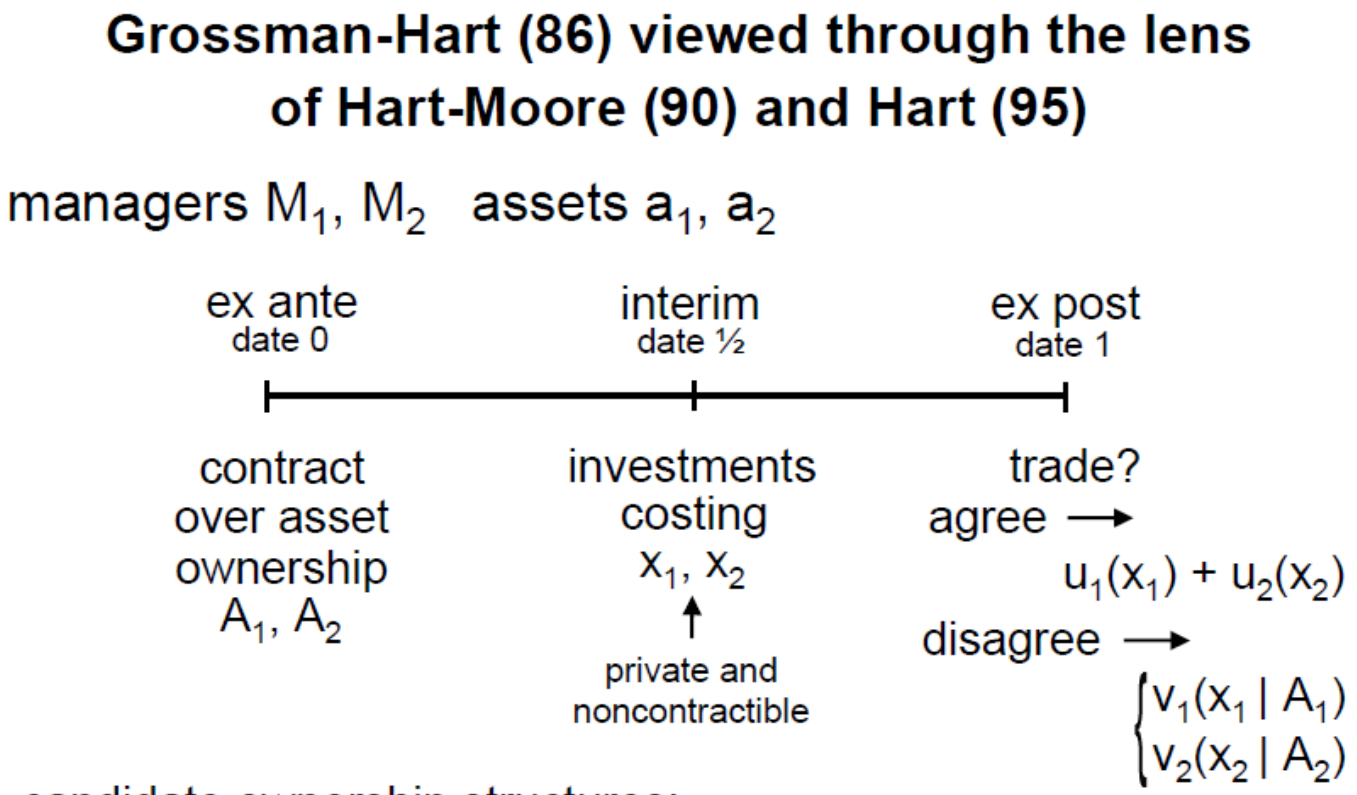
Grossman-Hart (86) viewed through the lens of Hart-Moore (90) and Hart (95)
managers M_1, M_2 assets a_1, a_2
ex ante
date 0
interim
date ½
ex post
date 1
contract over asset ownership A_1, A_2
investments costing x_1, x_2
private and noncontractible
trade?
agree
$u_1(x_1) + u_2(x_2)$
disagree
$v_1(x_1 \mid A_1)$
$v_2(x_2 \mid A_2)$
candidate ownership structures:
type 1 integration: $A_1 = \{a_1, a_2\}$ A_2 = null
nonintegration: $A_1 = \{a_1\}$ $A_2 = \{a_2\}$
type 2 integration: A_1 = null $A_2 = \{a_1, a_2\}$

Figure 1

A few assumptions, the main one at the end: Define ex post surplus as the sum of the u's minus the sum of the v's – that's what we're bargaining over at date 1 – and assume it is positive. Given the way I've set up the algebra, our private investments affect only our own utilities (on or off the equilibrium path), so there are no direct externalities. Another point to bear in mind – slightly more smuggled under the rug of algebra – is that the investments are in human capital only. Why? Because in the algebraic formulation, if the ex post bargaining breaks down, and you, for example, go off with your assets A_2, then those assets aren't affected by your (or my) investment. Let the u's and v's satisfy all the usual assumptions (smooth, strictly increasing and concave in investment).

Now for the main assumption: for i = 1 and 2,

$$\frac{\partial u_i(x_i)}{\partial x_i} > \frac{\partial v_i(x_i|\{a_1, a_2\})}{\partial x_i} \geq \frac{\partial v_i(x_i|\{a_i\})}{\partial x_i} \geq \frac{\partial v_i(x_i|\text{ null})}{\partial x_i} \geq 0.$$

This line-up of inequalities says that the marginal product of investment by agent M_i within the relationship (u_i) strictly exceeds all M_i's marginal products of investment outside the relationship (the v_i's). The second term in the line-up is the marginal product of investment if M_i owns both assets. From the second inequality we learn that this is larger than if M_i owns only asset a_i. The third inequality says the marginal product is least if M_i owns nothing. In other words, there are complementarities between access to physical assets and returns to investment in human capital. Think of this as an assumption over the sign of the second-order cross-partial derivative of returns with respect to physical and human capital.

We can view the v's as the managers' status-quo payoffs in an ex post bargain. Figure 2 illustrates.

FIGURE 2 near here

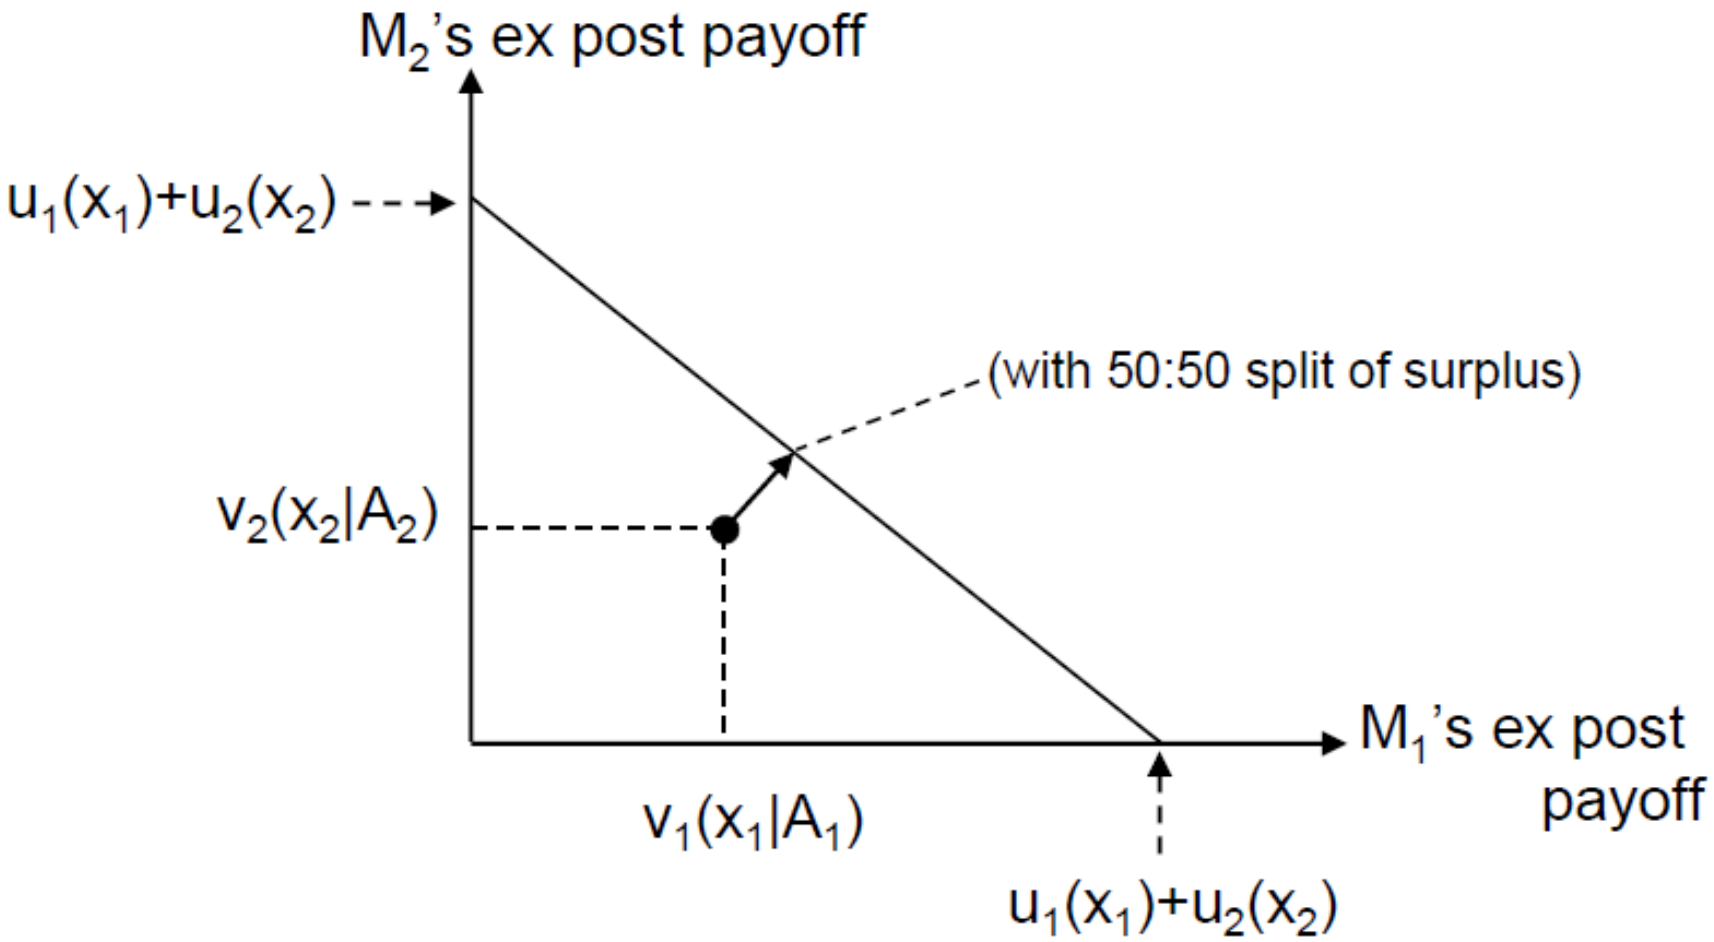
M2's ex post payoff
u1(x1)+u2(x2)
(with 50:50 split of surplus)
v2(x2|A2)
M1's ex post payoff
v1(x1|A1)
u1(x1)+u2(x2)

Figure 2

The 135 degree line is the date 1 utility frontier. The v's, the coordinates of the blob southwest of the utility frontier, is what happens if we don't reach agreement, the status quo. We are assumed to divide the surplus equally, taking us from the blob due northeast to hit the frontier.

Now the crucial question is: bearing in mind this 50:50 division of surplus (we have rational expectations about how the bargaining is going to work out at date 1), how will we choose, non-cooperatively, to invest at date ½? Here is the answer: for i = 1 and 2,

M_i chooses x_i to maximize

$$v_i(x_i|A_i) \;+\; ½\{\text{ex post surplus}\} \;-\; x_i$$

FOCs:

$$½\frac{\partial u_1(x_1)}{\partial x_1} \;+\; ½\frac{\partial v_1(x_1|A_1)}{\partial x_1} \;=\; 1$$

$$½\frac{\partial u_2(x_2)}{\partial x_2} \;+\; ½\frac{\partial v_2(x_2|A_2)}{\partial x_2} \;=\; 1$$

As agent M_1, I look ahead at my status quo payoff $v_1(x_1|A_1)$ plus my 50% share of the ex post surplus, minus my investment cost x_1 which I've assumed to be one-for-one.

From our respective first-order conditions (FOCs) for our investments, emerges the conclusion that both of us are underinvesting relative to first-best. In first-best we would equate the marginal product of investment to the marginal cost of investment (unity); and along the equilibrium path, the products of investment come wholly through the u's – the v's are irrelevant.

When we examine the private optimization problem, however, we discover that the weight we put on the first term on the left hand side of our FOCs is only ½, because of hold-up. Each of us anticipates that of every dollar we contribute to the joint surplus, fifty cents

are going to be stolen by the other party in the bargaining. Little wonder, then, that we have diluted incentives to invest. But there is an offsetting incentive: to feather our own nests. This is what the less obvious second term is all about.

I've got one eye looking towards what will happen on the equilibrium path (here the 50% dilution factor kicks in), and the other eye looking towards making sure I'll have a good outside option from where to start the bargaining. It turns out that I put a 50% weight on that too. The same applies to you.

So, instead of putting a weight of 100% on the social return, the u's, and nothing on the v's, we each put a weight of 50% on both. This ameliorates the dilution but doesn't offset it entirely because, from the line-up of inequalities, the second term on the left hand side of our FOCs is smaller than the first term, for each of us. Hence it doesn't add up to what a social planner would want: there is under-investment relative to first-best:

$$x_i < x_i^* \qquad \text{where} \quad \frac{\partial u_i(x_i^*)}{\partial x_i} = 1$$

After all this ground clearing, we at last arrive at Grossman and Hart's insight: Giving control of an asset to one agent improves his or her incentives to invest, but it reduces the other agent's incentive. See Figure 3.

FIGURE 3 near here

Imagine moving up a hierarchy, let's say from one extreme, type 2 integration, through non-integration, to the other extreme, type 1 integration. As we go up, I (agent 1) have ever greater incentives to invest. But you (agent 2) have ever smaller incentives. The proof comes from examining the pair of FOCs that constitute our Nash equilibrium at date ½. As the set of assets A_1 that I control expands, it pushes up the left hand side of the top equation. But it's a hard fact of life – of physics – that we can't both control the same assets. So

giving control over an asset to one agent improves his incentive to invest but reduces other agent's incentive:

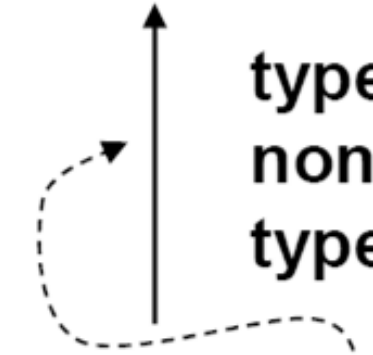

type 1 integration: $A_1 = \{a_1, a_2\}$ $A_2 = \text{null}$

nonintegration: $A_1 = \{a_1\}$ $A_2 = \{a_2\}$

type 2 integration: $A_1 = \text{null}$ $A_2 = \{a_1, a_2\}$

as we go up: $x_1 \uparrow$ but $x_2 \downarrow$

proof:

$$\tfrac{1}{2}\frac{\partial u_1(x_1)}{\partial x_1} + \tfrac{1}{2}\frac{\partial v_1(x_1 \mid A_1)}{\partial x_1} = 1$$

$$\tfrac{1}{2}\frac{\partial u_2(x_2)}{\partial x_2} + \tfrac{1}{2}\frac{\partial v_2(x_2 \mid A_2)}{\partial x_2} = 1$$

expands

contracts

Figure 3

though I'm getting an expanded set A_1 as we progress up the hierarchy, your set A_2 is shrinking. And that pushes down the left hand side of the bottom equation. There's a trade-off.

A gloss that you might put on this theory, due I believe to Bengt Holmstrom, is that assets correspond to bargaining chips. Bargaining power emanates from control over assets. And bargaining power is a scarce commodity. It would of course be wonderful if we could both have control over both assets. But that's infeasible, and we must allocate the chips judiciously.

Let me make a few remarks about the strengths and weaknesses of the Grossman and Hart model. Its strength is that it delivers a great insight. That insight is also a simple insight. Of course it is, because all great insights are simple. (Dispiritingly, not all simple insights are great.)

Grossman-Hart (1986) introduces into the formal literature the notions of power and control. It's obvious, isn't it, that we need to have power and control in our understanding of society – we need to have those as ingredients in our modelling? Yet the economist's bedrock, the Arrow-Debreu framework – and even an embellished Arrow-Debreu framework encompassing all of agency theory and information economics – does not admit power and control, a rather naked force that seems extremely important.

Within a world of incomplete contracts, Grossman and Hart proposed equating power with control over *non-human* assets. In so doing, they decisively moved away from the earlier literature dating back to Coase (1937) that had – arguably – gotten mired in a discussion of human capital and an idea that Coase had put on the table: that if I work for you, I must do what you tell me to do; whereas if I am an arm's length contractor, I needn't. That loose idea sat in the literature until 1986 when Grossman and Hart shifted it entirely. They argued that we shouldn't think in terms of someone having direct control over other people, because a person's human capital is inalienably his or her own (I've got mine, you've got yours, and that's all there is to it). Simply relabeling someone an employee cannot solve agency problems. Instead, Grossman and Hart argued, we should think about control over non-human assets. Only *indirectly* might that form of control give someone control over other people.

It's clear that in practice, control over non-human assets matters a great deal. In the absence of any specific agreement about how a particular asset is to be used, the party who owns the asset has control over it: the right to decide how to use it, who to give access. But such power is a scarce commodity and needs to be allocated across society, asset by asset. This allocation is done by means of contracts. Although Grossman and Hart's is a model of power in society, their society is legally structured. Their equating of ownership with control conforms well with the views of lawyers.

Rather unusually, the model has only one force at work. It happens to cut two ways. If you've got control, then I haven't; and vice versa. In other words, the cost of integration is simply the dark side of the benefit. That elegant idea wasn't in the earlier literature on the firm. In the earlier literature, because there was a strong sense that integration, the creation of a single firm, reduces agency problems, there were many apparent benefits. Costs were of another variety altogether, perhaps something to do with limited managerial span. As a general rule of research, it's a pity to write down a model in which benefits and costs are distant cousins. Such a model can seem too kitchen sink. An ideal theory has a single force working in two ways, so that the costs and benefits are two sides of the same coin. The Grossman-Hart (1986) paper has exactly this feature. Their model is beautiful in its simplicity; and from the simplicity comes a broad set of applications.

(This insight – that the cost of integration is the obverse of the benefit – is unique to Grossman and Hart (1986). It would be a mistake to think that all Grossman and Hart did was to put into maths what was already in the literature. This is not a matter of maths versus no maths; it's a matter of clear thinking versus less-than-clear thinking. A precise contribution of Grossman and Hart's 1986 article was to show that the costs and the benefits of integration stem from the same force.)

Not only is the model beautiful, it also makes very sensible predictions, such as: (1) Important people should be bosses, i.e. should be owners. (2) Complementary assets should be owned together. (3) Assets that have nothing to do with each other shouldn't be owned together. If (2) and (3) are combined, we arrive at Dennis Robertson's analogy, quoted by Coase (1937), of "lumps of butter coagulating in a pail of buttermilk": firms are the lumps of butter in the market economy, the buttermilk. Or, to put it less poetically, the optimum size of the firm is interior: it's not zero, but it's not the whole economy either.

When teaching this paper, I find this last point strikes home the most. Grossman and Hart (1986) is as much a theory of the market, of the benefits of decentralization, as it is a theory of the firm. Although these are equivalent ways of looking at the paper, the interpretative stress is different. Earlier in Figure 3, when moving up the hierarchy of ownership structures from type 2 integration through to type 1, it was easy to skip over the middle option: non-integration. (In *A Fish Called Wanda*, Kevin Kline stole the show – and won an Oscar – playing nasty-but-dim Otto. The running gag had Otto replying to any triplet of choices, "What was the middle one?") The fact that non-integration is typically superior to the other two ownership structures is a terrific argument in favour of the market economy and decentralization.

The theory has profoundly influenced many fields. I'm particularly keen on the application to corporate finance, with Aghion and Bolton (1992) the starting point. And Grossman-Hart (1986) will find ever more application to political economy: How can we model government without modelling power and control, and how can we understand power and control without their idea?

Weaknesses. Etiquette perhaps demands that I not write about these in a celebratory volume, but one weakness of the model just presented is that, in order to yield most of the sensible predictions, we need to make assumptions signing second-order cross-partial derivatives. It seems a pity that an idea so apparently clear-cut and intuitive descends – should that be ascends? – into a discussion about second derivatives. Somehow, I fear we've climbed a derivative unnecessarily. One ought to be able to tell this story about power and control at the level of first derivatives (or below) and not have to worry about the signs of second-order cross-partials.

Another worrying sign – thanks here to de Meza and Lockwood (1998) – is that the predictions are vulnerable to putting a different bargaining game into the model. Instead of having a 50:50 division of surplus from a status quo pair of payoffs, if an outside option game of bargaining is used then de Meza and Lockwood show that the predictions of the model can change radically. I think this may be connected to my first point: we're too many derivatives up for comfort.

The Grossman-Hart model perhaps puts too great an emphasis on interim inefficiency. Hold-up is important, but looking round the world it seems that ex post inefficiencies are

even more important. Coase (1960) is acting as a straitjacket to ensure that ex post bargaining is fully efficient. Although a world of frictionless trading at date 1 is a sensible place to begin, I'd prefer to get away from full ex post efficiency. Also, it would be good to jettison that middle date ½ and investments: they snarl up the analysis and force us to think about second derivatives. I'm not the first to say these things. Oliver and I have worked on the matter more recently, making use of the idea that contracts may act as reference points (Hart-Moore, 2007, 2008); and he's taken this on further in his own work (Hart, 2008, 2009, 2013).

Finally by way of criticism, the elephant in the room: Why are contracts incomplete? I don't see this as a fatal criticism of the Grossman and Hart paper. In 1986 they took the bold step: Let's *assume* contracts are incomplete – there are some wonderful footnotes in the paper wrestling with ideas to justify this assumption – and let's put power and control into the picture. But haven't they had a hard time from some pure theorists! I don't think it's Grossman and Hart's fault. It's our fault. Mea culpa. We've had twenty five years to come up with a watertight theory of contractual incompleteness and we haven't succeeded yet. Eventually someone will.

The model I first heard Oliver give at the LSE wasn't the one I've described above. It was another, very different, model – published later in a book edited by Razin and Sadka. I'd love to have been a fly on the wall when Sandy and Oliver were doing their research. My hunch is that they began by writing down this other model. Then, to deal with the constraints imposed by Coase (1960), the model had to be embellished with investments and hold-up to generate inefficiencies. Let me describe the other model (see Grossman and Hart, 1987).

Take a buyer and a seller, and dispense with date ½ and investments. The buyer's value v and the seller's cost c are independently and uniformly distributed ex ante. First-best has ex post trade if and only if v exceeds c. But, crucially, there is no ex post renegotiation. Contracts, written at date 0, are assumed to be wildly incomplete, with no attempt at building in mechanisms to make the trading price contingent on revelations of the state of nature. Simply assume that a contract specifies a trade price p_1 and a no-trade price p_0. (The paper has some interesting early rationalisations for why this may not be such a bad assumption.)

The contract also specifies which party, if either, has control. In a situation of non-integration, trade has to be voluntary, so the gap between the prices, $p_1 - p_0$, has to be less

than v and more than c in order to induce both parties to trade. Or the buyer might be given control, and then he'll have the ability to force trade through even if the seller doesn't want to: there will be trade whenever v is greater than $p_1 - p_0$. Or the seller might be given control and there will be trade whenever c is less than $p_1 - p_0$.

In sum, the design of a contract involves: first, deciding which party, if either, has control; and, second, choosing the values of p_1 and p_0. The analysis boils down to studying diagrams like those in Figure 4.

FIGURE 4 near here

Remember the distributions of v and c are uniform and therefore lie in a rectangular box in v-c space. In first-best, we want trade if and only if we are north west of the 45 degree line, where $v > c$. But we're limited by the crudeness of the contracts that we have at our disposal. Under non-integration, there is trade if and only if $v > p_1 - p_0 > c$; see the top diagram. Trade happens in the shaded area. There is undertrading – too little trading relative to first-best. However, given the position of the box, it turns out that both of the other two control structures (buyer control, seller control) are worse.

If the box is in a different place, as in the bottom left diagram, it turns out that non-integration and seller control are both beaten by buyer control. With buyer control, there is trade if and only if $v > p_1 - p_0$. Notice that we've got a lot more trading here, but some of it is inefficient: we have overtrading as well as undertrading.

In the bottom right diagram, we have a yet differently placed box (the dimensions and positions of these boxes are worth thinking about), and now it turns out that seller control is the optimal contractual arrangement.

What I really like about this simple model is that it has nothing to do with second derivatives, or even first-order conditions. The model has a naked ex post inefficiency – of course we should be anxious about Coase (1960) – but it distils the essence of the matter. Contracts are incomplete. Someone has to make a decision. Inevitably, the other party may find that decision costly, and thus there is trade-off between costs and benefits. The three

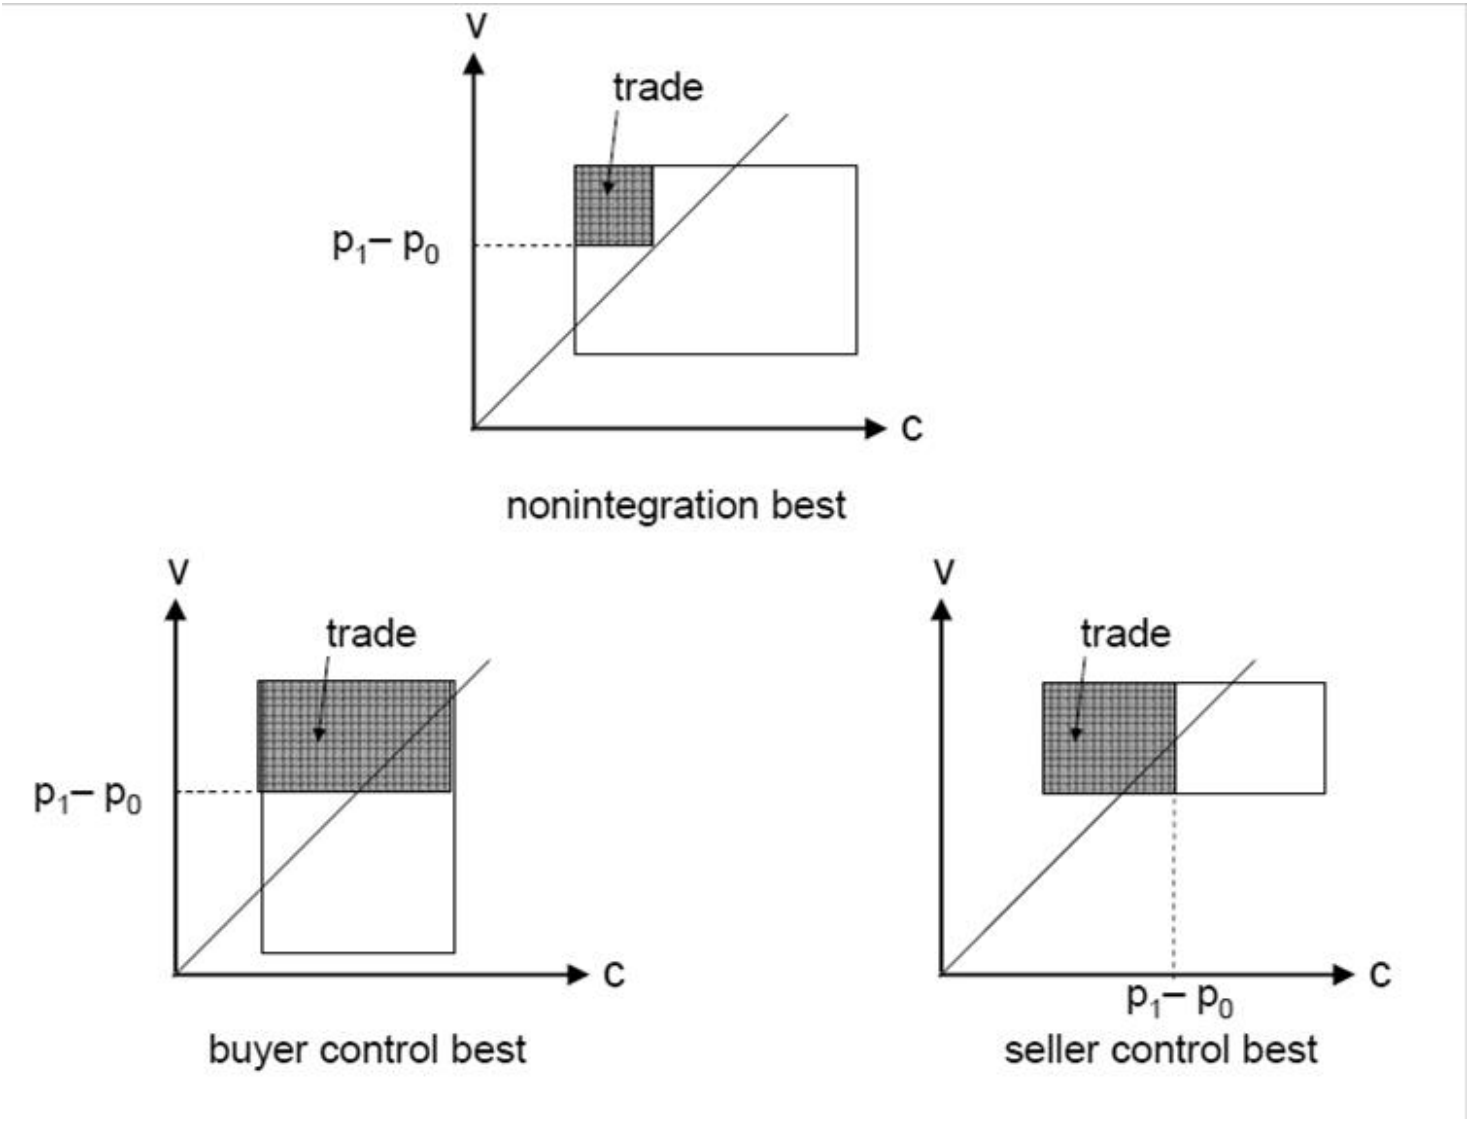
V
trade
$p_1 - p_0$
c
nonintegration best
V
trade
$p_1 - p_0$
c
buyer control best
V
trade
c
$p_1 - p_0$
seller control best

Figure 4

diagrams capture that idea beautifully. I believe this may be the right line for future research: the role of a contract is to steer parties close to their ex post efficient frontier. Yes, they may be able to rely on renegotiation at date 1 to take them even closer to that frontier, but there are sensible reasons why they may have difficulty reaching it. I think this much cruder way of capturing the effects of allocating control power has a lot to recommend it compared to the rather subtle business of trying to induce interim investments.

Sandy and Oliver thought up their theory in a week, wrestling with it together in Chicago, provoked by the AT&T case. (Incidentally, I'm not sure how much of this is true; but it is the way I remember it.) The next bit I do remember well. Oliver flew directly from Chicago to Penn for a summer workshop where I was too. He was palpably excited, saying he'd just had the best week of his life with Sandy, and that they had perhaps come up with something important. As he told me, he was trembling in a rather un-Oliver way, even in the sweltering heat as we devoured hoagies from the vans that park just off Locust Walk. Oliver said that he thought their idea was their best thing yet: he felt that nothing they had written before came close. This was hot off the press, the day he'd just left Sandy. I thought, wow, this from two guys who had, as sole authors, already revolutionised asset price theory (in the case of Sandy) and general equilibrium theory (in the case of Oliver). I was caught up in Oliver's excitement and I told my friends on the quiet that They Were On To Something Big. Of course, I didn't know what it was.

It's worth saying that when the paper did finally emerge, not everyone fell about declaring it a masterpiece. One LSE professor, who will remain nameless, said: "The theory is not very interesting. It'll never go anywhere. But the application to the insurance industry in the last section – that might be worked up a little. That's the only part that will be remembered."

When Kydland and Prescott's real business cycle theory was first presented, it fell to Ken Rogoff to act as discussant. Rogoff pointed out that there were two tests of a true scientific revolution. At first the idea must seem mad. Then it must seem obvious. And he felt that the RBC theory certainly passed the first test. Now Grossman and Hart's 1986 paper was also a scientific revolution – to my mind just as important to microeconomics as RBC theory is to macro, because, as I've emphasised, their theory puts the idea of power and control into mainstream economics, and without these we can have little handle on many important economic phenomena. Ironically, maybe the greatest legacy of Grossman and Hart (1986)

will be outside the relatively narrow confines of organizational theory. But arguably Grossman and Hart's theory fails Rogoff's first test, because it has never seemed mad. Their theory seems almost self-evident, which is perhaps the greatest of compliments.

Oddly enough, although the property rights theory has become an indispensible tool for many applied economists, among pure theorists it has proved harder to win friends, largely because we still lack a theory of contractual incompleteness. Martin Hellwig, hearing of the title to Jean Tirole's 1994 Walras-Bowley Lecture (published in *Econometrica*, 1999), ribbed "Incomplete Contracts: Where Do We Swim?"

As has been said of the decentralised competitive system, it works in practice, but does it work in theory? Well, the authors in this book will be discussing these and many other fascinating matters, all because 25 years ago these two gentlemen had a stroke of genius.

References

Aghion, P. and P. Bolton (1992), "An Incomplete Contracts Approach to Financial Contracting", *Review of Economic Studies*, Vol. 59, pp. 473-494.

Coase, R. (1937), "The Nature of the Firm", *Economica*, Vol. 4, pp. 386-405.

Coase, R. (1960), "The Problem of Social Cost", *Journal of Law and Economics*, Vol. 3, pp. 1-44.

de Meza, D. and B. Lockwood (1998), "Does Asset Ownership Always Motivate Managers? Outside Options and the Property Rights Theory of the Firm", *Quarterly Journal of Economics*, Vol. 113, pp. 361-386.

Grossman, S. and O. Hart (1986), "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration", *Journal of Political Economy*, Vol. 94, pp. 691-719.

Grossman, S. and O. Hart (1987), "Vertical Integration and the Distribution of Property Rights", in A. Razin and E. Sadka (eds.) *Economic Policy in Theory and Practice*, London: McMillan, pp. 504-546.

Hart, O. (1995), *Firms, Contracts, and Financial Structure*. Oxford: Oxford niversity Press.

Hart, O. (2008), "Reference Points and the Theory of the Firm", *Economica,* Vol. 75, pp. 404-411.

Hart, O. (2009), "Hold-Up, Asset Ownership, and Reference Points", *Quarterly Journal of Economics,* Vol. 124, pp. 267-300.

Hart, O. (2013), "Noncontractible Investments and Reference Points", *Games,* Vol. 4, pp. 437-456.

Hart, O. and J. Moore (1990), "Property Rights and the Nature of the Firm", *Journal of Political Economy*, Vol. 98, pp. 1119-1158.

Hart, O. and J. Moore (2007), "Incomplete Contracts and Ownership: Some New Thoughts", *American Economic Review,* Vol. 97, pp. 182-186.

Hart, O. and J. Moore (2008), "Contracts as Reference Points", *Quarterly Journal of Economics,* Vol. 123, pp. 1-48.

Tirole (1999), "Incomplete Contracts: Where Do We Stand?", *Econometrica,* Vol. 67, pp. 741-781.