

Geoffrey Brennan

The Division of Epistemic Labour*

Abstract: The paper mobilizes Adam Smith's treatment of the division of labour in relation to the production, consumption and exchange of knowledge. One aspect of this mobilization deals with the epistemic demands that exchange makes on its participants. The other deals with increasing returns in the provision of knowledge itself, treating knowledge creation as just another example of specialization and exchange. These two aspects come together in relation to the epistemic demands associated with assessing knowledge quality. These demands differ according to whether the knowledge is embodied in products or whether the knowledge is an object for its own sake. It is argued that disciplines play a critical role as institutions for meeting the epistemic demands that the division of labour creates in the 'knowledge' case.

1. Throat-clearing

In this paper, I want to explore the idea of the division of epistemic labour—and what is required to make it work tolerably well.

It is natural for an economist (as I more or less am) to begin that exercise by 'going back to Adam Smith'—partly because the division of labour plays such a central role in Smith's basic picture of the economic order; partly because Smith quite explicitly identifies the division of epistemic labour as one aspect of his overall account; and partly because aspects of Smith's account are quite distinctive and this distinctiveness seems to have been largely overlooked in subsequent literature.

Although my respect for Smith is considerable, I think some aspects of his account of the division of epistemic labour are questionable, or partial. Simply put, I think Smith attends to only part of the knowledge economy—and that the part that operates outside the narrowly commercial sphere has some distinctive properties to which Smith did not attend and require institutional supports that Smith consequently does not consider. The road plan is as follows.

In *section 2*, I examine Smith's general account of the division of labour, emphasizing its distinctive features. In *section 3*, I draw attention to an epistemic challenge to the division of labour for all goods—a challenge which Smith back-

* My thanks are due to Michael Baurmann for extremely helpful editorial comments and to Nic Southwood, who acted as commentator on this paper in its originally presented form. Standard caveats apply.

grounds but which becomes central in dealing with the case of ‘knowledge’ as a good. *Section 4* directs attention to the division of labour as applied expressly to the production and diffusion/consumption(?) of knowledge as one specific good; and how the general epistemic challenge discussed in *section 3* manifests itself in the case of knowledge. *Section 5* concerns the special role that disciplines play as institutions that meet that epistemic challenge, and a little about the ‘costs’ as well as the benefits that disciplines offer in this regard. *Section 6* offers a summary conclusion.

2. The Division of Labour in the Economy of Goods

“The greatest improvement in the productive powers of labour, and the greater part of the skill, dexterity, and judgment with which it is any where directed, or applied, seem to have been the effects of the division of labour.” (Adam Smith WN I.1.i)

So Adam Smith begins his classic treatise on the economy.

In making this remark, Smith did not, I think, see himself as saying anything especially unfamiliar. The idea of the division of labour had been abroad more or less since Plato’s time; and the details of Smith’s striking (though arguably somewhat misleading)¹ example of the pin factory were taken directly from William Petty. Equally, of course, nowadays, the idea of specialization and the division of labour are, as Schumpeter notoriously put it, “clichés” in economic understanding. Smith seems, in that sense, to be just another patron in the history of that idea.

It is therefore worth emphasizing that Smith’s understanding and use of the ‘division of labour’ concept is distinctive. In other classic accounts (Plato, Aristotle, St Paul, and for that matter in the familiar Ricardian account of ‘comparative advantage’) the grounds for specialization lie in the natural differences between individuals (or as in the Ricardian case, regions)—so that specialization is a matter of developing the distinctive talents with which individuals² are by nature endowed. By contrast, for Smith, the division of labour (specialization) does not depend at all on any such natural differences—and could in principle arise in a world of naturally identical persons/regions. In Smith’s account, specialization is profitable because of what economists would now call increasing returns. Individuals develop their skills by repetition and dedicated attention—so that productivity increases when individuals specialize, even when those individuals are by nature identical (as Smith believes they more or less are!)

One significant implication of this fact is that, whereas in the rival accounts, the extent of specialization is limited by the natural differences in individuals’

¹ Buchanan for example suggests that the ‘economies of scale’ that are generated by specialization are ‘external’ to any firm or industry—and criticizes the pin factory example specifically on this account.

² Or in the Ricardian case, the distinctive natural properties of regions (climate, fertility, etc).

productive capacities (or regions), in Smith the scope for an extended division of labour is limited only “by the size of the market”—as Smith puts it in the title to chapter 3 of Book I. So the division of labour, on the Smithian picture, is virtually *unlimited*: as population increases, specialization becomes more and more fine-grained (including specialization in the acquisition and development of knowledge—an aspect of the division of labour that Smith explicitly mentions in chapter 1, and which we shall take up in greater detail below).

The ‘increasing returns’ aspect of Smith’s account has been largely lost to subsequent writers—Schumpeter’s relegation of specialization to the status of cliché notwithstanding. Smith shares some of the blame for this loss. In the later development of his analysis of the ‘stationary state’ (in WN book II), there is no mention of increasing returns; and so Samuelson can re-construct the Smithian general equilibrium model and resurrect Smith as an ‘economic theorist’ without increasing returns playing any role. Ricardo’s model of the gains from trade is basically a ‘*constant* returns to scale’ model—and the standard general equilibrium in mainstream economics (the Arrow-Debreu model) does not depart from Ricardo in this respect. The appeal to constant returns to scale is perhaps understandable, because the assumption of constant returns greatly simplifies the analytics. Nevertheless, the increasing returns aspect of Smith’s picture is critical to his own account of the ‘wealth of nations’: and for that matter, to explaining what we have in fact observed in the West over the last three hundred years—namely, rapidly increasing aggregate population *alongside* rapidly increasing real per capita income.³

Of course, Smith does not need to *deny* natural differences in order to establish the increasing returns aspect of the division of labour. He himself uses the notion of comparative advantage at various places—and even makes appeal to the phenomenon of diminishing returns in particular cases and specific industries.⁴ And he does not claim that increasing returns applies everywhere. In particular, he thinks it does not apply in agriculture—a belief that led him to the prediction that agricultural prices would increase over the long term.⁵ Needless to say, he seems to have been quite wrong on this latter point: agricultural prices have been if anything in long term secular decline through most of the last century or so. Nevertheless, Smith’s general picture is of a trading nexus made up of naturally identical individuals, all participating in a dynamic process characterized by ever-increasing specialization as the numbers of traders increases. The result is a ‘general opulence’ that diffuses itself among an expanding population; and provides the grounds for increased population and thereby for yet further refined specialization (and yet greater opulence).

It may be worth emphasizing the role that exchange plays in this picture. On Smith’s account, exchange is significant because it permits specialization in production to take place. He recognizes, of course, that exchange itself is based

³ A trajectory that the Malthusian ‘diminishing returns’ model of population size seems to render almost inconceivable.

⁴ For example, he notes in relation to mining that, as the line of lode runs deeper underground, the extraction of ore becomes more and more expensive.

⁵ Along lines more recently associated with Raoul Prebisch (1950), and William Baumol’s (1967) “economics of unbalanced growth”.

on the mutual benefit that parties derive from it. But the gains from exchange as such are not the primary ingredient in his normative defense of the market order.

To make the relevant distinction clear, it may be useful to offer a simple example. In a paper that is sometimes assigned to undergraduates, Radford (an Australian, hence the reference!) describes the emergence, in a prisoner-of-war camp, of prices for the various items included in the Red Cross packages that would periodically arrive for the inmates. Differences in tastes for the various objects, differences in attitudes to risk (since the length of time before the next package arrived was unknown), differences in discount rates, were sufficient to constitute a basis for exchange among the prisoners. Trade occurred and prices emerge. And the parties are made better off by such exchanges. But *these* gains are, by Smithian lights, entirely second order. Compared with the increased productivity accruing from specialization (of the order of perhaps 4800 fold in the ‘trivial’ case of pin manufacture), the gains from exchange *as such* are to be seen as very small beer. Trade is crucial because it permits the separation of production and consumption without which specialization in production can never get off the ground. But the primary source of mutual gains lies, as we might put it, in technology rather than in catallaxy as such.

3. The Epistemic Demands of the Division of Labour

Although it is not an aspect of the division of labour that Smith emphasizes, it is worth noting immediately that the economy that Smith describes (and extols) makes significant epistemic demands on its participants. Consider a simple example. In a fascinating recent book, Bee Wilson (2008) recounts the story of how food ‘additives’ constituted a mounting problem through the nineteenth century as the division of labour proceeded. In a less ‘divided’ time when people ate exclusively things that they themselves had grown, killed and/or cooked, they could know with a fair degree of certainty the edibility of that diet. Provided they were not too experimental in the mushrooms they collected or in their meat-storage practices, they could be reasonably assured that the food they ate would not poison them. But as a larger and larger proportion of their diet became composed of items grown, prepared, and cooked by others—and specifically the anonymous others that made up the great company of traders in the market-place (both producers and intermediate traders)—knowledge of what exactly they were eating and how it was prepared became less a matter of personal control and more a matter to be taken on trust. The question of whether that trust could be justified became accordingly an increasingly central issue for the emerging trading economy—and not just in relation to food. One could be reasonably sure of the quality of the elements that went into the building of one’s house if one constructed it oneself, even if the workmanship was somewhat rough and ready: but a house constructed by others had to be carefully inspected. One’s own horse was one whose reliability was, at least,

something you could presume familiarity with: the hired hack was an unknown quantity. And so on.

Trade involves trust; and trust can be abused. As the trading nexus expanded beyond the local village to the larger town and eventually to the big city, the natural disciplines of local reputation became increasingly strained. A *local* baker or butcher or brewer was liable to personal abuse and significant loss of reputation if his product proved unsatisfactory. Dissatisfied customers could spread the word about the shoddy bread or the rotten meat and throw stones on his roof to express their dissatisfaction. But the prepared sausage, the exotic cake and the bottled beer brought in from the London specialists could well have ingredients of unknown origin and dubious effect. And there was little the disappointed individual customer could do to damage either the maker's reputation or his person!

The simple general point is that every act of trade creates a 'principal/agent' problem. When the principal and the agent are the same person, natural self-interest will ensure that the agent's motives are directed at the best interests of the 'principal'. In all other cases, this is a contingent matter. And if the division of labour is to flourish and grow, it is a matter that has to be somehow resolved.

Smith might, of course, point to the brute fact of flourishing markets in his own day to argue that this is an issue that *was* solved—somehow. That is, he could reasonably argue that the fact that we look to the butcher and baker and brewer for our dinner is itself testimony that by and large we trust those agents to provide us with better meat, bread and beer more cheaply than we could make for ourselves (or acquire elsewhere). But Smith's claim that the division of labour is limited by the *size of the market* (understood as the population of potential traders) is at best only partial and at worst somewhat question-begging. If the division of labour is not to be limited instead by the extent of trustworthiness, then the principal-agent problem that dogs all exchange must somehow be solved. And providing some account of what that solution might be, and whether better and worse solutions are available, seems to be a critical element in any account of the 'wealth of nations' in which exchange and specialization are central.

A couple of remarks about trust/trustworthiness may be in order here.

First, in the rationalist account of the phenomenon, the chief driving element is not trust but trustworthiness: trust is seen to follow in the wake of trustworthiness as a rational response. Rational agents trust those who are trustworthy. But this possibility depends on the capacity of agents to distinguish trustworthy from untrustworthy potential partners. Clearly then, in such a setting, it is not enough for a potential trading partner to *be* trustworthy—she must be *known to be* trustworthy. For this reason, 'signaling' trustworthiness in a reliable manner becomes an essential feature of the trust game.

Second, if acquiring knowledge as to who is trustworthy and who not, is costly, then it may pay individuals to 'trust' to the market rather than acquire the relevant information directly. The thought is that if enough *others* have checked the trustworthiness of a given trader, then the fact that that trader is still in business is some evidence that she has survived the test. So each can effectively free-ride on the information-acquisition activities and/or experience

of others. However, that free-riding possibility creates a niche for untrustworthy players. The free market may well exhibit the feature that individual traders have less information about other traders than is optimal. So there will be more fraud than is optimal.

It will be clear that the epistemic problems associated with exchange are larger the more costly it is to ascertain product quality. Some goods are such that their quality is readily accessible via inspection—or at least, quality is assessable at low cost. Other goods are not. In the latter case, the provision of relevant information may be a specialized activity within the market. Some car mechanics specialize in doing inspections of second-hand cars; the Michelin guide provides information about the quality of restaurants; credit rating agencies (like Standard & Poors, and Moody's)⁶ purport to provide information about the credit-worthiness of different financial assets; Robert Parker makes a steady income from providing information about the quality of specific wines and of different vintages in various locations; and so on. These market institutions for the provision of information pose, of course, the same trustworthiness and signaling issues as do the manufacturers/providers of the products that the information-providers assess. But it is interesting to note that the division of labour allows specialization in the particular business of providing reliable information and signaling trustworthiness. In this instance at least, demand goes some way to creating its own supply.

There is a vast and complicated literature on the signaling of trustworthiness and on institutional devices that might support incentives to behave in a trustworthy manner. I do not intend to explore that literature here (except in one particular case, of which more anon). I do however want to underline the role that reputation plays in markets where product quality is uncertain. In some markets, the image of entirely anonymous traders makes perfect sense: these are markets where epistemic problems are second-order. But in many markets, traders are anything but anonymous: indeed, their 'name' is, if not 'everything', certainly a very great deal. Traders will be careful to develop and preserve their reputations—and will rationally promote them, once acquired. This is not an incidental feature of markets: it is a central enabling feature in any market where product quality is less than completely self-evident.

4. The Division of Labour and the Production and Consumption of 'Knowledge'

So far, I have attended to the epistemic demands of exchange in relation to ordinary goods—and the role that trustworthiness, and institutions that promote it, might play in supporting the operation of ordinary markets. But this is to attend to only one aspect of the knowledge economy—knowledge about goods

⁶ As the recent financial crisis has indicated in relation to the prominent credit-rating agencies, that specialization can just serve to magnify the trust problem—though of course, one dubious record does not establish that the whole enterprise of specialization in the provision of market-based knowledge is suspect.

and traders. It is now time to attend to the production and consumption of knowledge more generally, as a good. And again, we can begin with Smith:

“In the progress of society, philosophy or speculation becomes, like every other employment, the principal or sole trade and occupation of a particular class of citizens. Like every other employment too, it is subdivided into a great number of different branches, each of which affords occupation to a peculiar tribe or class of philosophers; and this subdivision of employment in philosophy, as well as in every other business, improves dexterity, and saves time. Each individual becomes more expert in his own peculiar branch, more work is done upon the whole, and the quantity of science is considerably increased by it.” (WN I.1.9)

Smith makes this comment in the context of a discussion of “machines”—which are of course an important bearer of the “*improvement in the productive capacity of labour*”. No surprises here. The role of the accumulation of science in human advance is a familiar theme. But what Smith emphasizes is not so much the role in human progress of knowledge *accumulation* over time—associated perhaps with the picture of each generation standing on the shoulders of its predecessors—but rather the idea that the pursuit of knowledge is itself a feature of, and internal to, the division of labour.

If we think of knowledge more as a final product than as an input into productive activity, we can treat the organization of enquiry as more or less analogous to the organization of production. Or at least we can, if we can maintain a plausible distinction between the production and consumption of knowledge and an epistemic analogue to ‘exchange’. And here there is I think an interesting issue.

In the cabbage market there are producers and consumers of cabbages and the two classes are readily identified. In the knowledge market, the distinction is less clear. When a reliable expert engages in some exercise in enquiry and ‘produces’ some knowledge, it can be transferred to others—via a process that might be thought of as somewhat analogous to exchange.⁷ If you are a ‘consumer’ (a down-stream recipient) of that knowledge, you come to ‘know’ the results of the

⁷ Although I think the distinction between consumers and producers in the knowledge context is clear enough, the characterization of the relation between them as ‘exchange’ is more troublesome. In some cases (cases of directly commercial use, say) it may well be that giving you the use of information/knowledge diminishes its value to me—perhaps destroys its value entirely—or that the ‘terms of knowledge exchange’ preclude the producer from using the knowledge in the way the consumer intends to use it. But even where this is so, the producer does not cease to know what he knows. When A and B exchange oranges and apples, A gives up oranges to get access to B’s apples, which B gives up. When A and B ‘exchange’ knowledge, both retains the knowledge they previously had (though as I say there may be limits on how they can use the knowledge they possess). To the extent, however, that knowledge is valued for its own sake, each does not lose the full value of his own knowledge by sharing it with the other. This fact relates to the ‘public goods’ nature of at least some knowledge. The point is that to characterize the relation between producer and consumer of knowledge as an exchange relation rather strains the notion of exchange as Smith (and his modern day fellow-travelers) typically think of it.

experiment (or whatever). But it is clear that you yourself do not possess direct epistemic warrant for the validity of the relevant claims. You ‘know’—to the extent that you do know—by virtue of the authoritative procedures of knowledge acquisition/production being faithfully followed by others. You ‘know’, not by experimentation or verification, but by virtue of the testimony of the original producers.

You could, of course, be a ‘consumer’ of knowledge in a more indirect sense. To take an example familiar to economists, it is frequently pointed out to first-year students that no-one in the world knows how to make a pencil. One person knows how to fell the trees; another how to shape the wood; another how to make the carbon centre (the ‘lead’); another how to make the glue that sticks the whole thing together. But no-one knows how to do all those things. And it is a standard claim about the decentralized properties of market order that it is not necessary that anyone knows. Provided those that (know how to) do the various bits in the production process are appropriately coordinated via the market, the various bits of knowledge at stake are combined to produce pencils of quality via a long chain of cooperation in which no single person has knowledge of all the elements. When each person in that chain purchases the partially-made pencil from the previous point in the chain, she acquires the fruits of the knowledge embodied in the pencil to that point. In this kind of process, however, we do not say that she acquires the knowledge as such: she simply acquires the fruits of that knowledge. The whole point of the example is that she does *not* acquire that knowledge as such—and she does not need to.

On this basis, it is appropriate to make a distinction here between ‘embodied’ and ‘disembodied’ knowledge. Knowledge that is embodied in products can be ‘passed on’ in the products in which they are embodied. Knowledge that is disembodied—knowledge valued for its own sake—cannot be. Disembodied knowledge must be communicated directly and possessed in the same *form* as the originator possesses it. It must be ‘shared’.

This sharing requirement seems to me to limit the advantages of the division of labour in the disembodied knowledge case. Consider for example a community in which each person knows something different— K_1 is the set of propositions known by individual one; K_2 the set known by individual two; and so on. Suppose these sets are all different. Each derives some utility let us suppose from that knowledge just by virtue of the knowing—each takes pleasure in knowing whatever it is that she knows. It satisfies some basic human propensity towards intellectual curiosity. Now, we might plausibly say that the community knows a great deal more than any one individual— $n \cdot K_i$ propositions. But that is a total derived from simple aggregation; it does not add to the utility derived by anyone in the community. That utility remains simply a function of the amount that each knows. To the extent however that that knowledge can be embodied in products (or part-products) that can be exchanged, then everyone in the community can benefit from the knowledge that others have—each benefits from the aggregate knowledge of the society. In the embodied case, it is important to the benefits which each derives that the knowledge possessed by others is different. In the disembodied case, this does not matter. Analogously, in the embodied

case, doubling the number of agents makes a difference to the benefits that each derives; in the embodied case, doubling the number of agents doesn't make any difference to the amount that any one of them knows and hence cannot increase average benefit. At first cut, increasing returns associated with specialization is restricted to the embodied case.

In fact, I think that first cut claim is a slight overstatement. I think there are *some* 'increasing returns' in the disembodied knowledge case. But the claim that increasing returns are much more significant in the embodied case strikes me as generally valid and important. With knowledge that is valued for its own sake there is a specialized role for correcting mistakes (since I am taking it that in general the demand is for propositions that are 'true') and for discovering more 'interesting' facts/ideas—but the fact that in the disembodied case, I do not benefit significantly if the things that I 'know' are distinct from the things that you know is a clear restriction on the desirability of specialization.

There is a further difference between embodied and disembodied knowledge relating to the fact that the epistemic demands involved in assessing the quality of the 'product' (demands of the kind explored in the previous section) are different in the two cases. Any knowledge that is embodied in product can be assessed for its validity indirectly—by appeal to the quality of the product. So, I do not need to know anything about how the internal combustion engine works (I don't) in order to be able to tell which of two cars 'drives better'. I do not need to know the slightest thing about the operation of computers to know which delivers my email faster and which software I find more user-friendly. I do not need to know and I do not need to ask!

Of course, assessment of product quality might nevertheless require a certain expertise. Consider an example. Suppose, for the purposes of the argument, that 'economic knowledge' is valued exclusively for its capacity to deliver better economic outcomes. And since we all bear the effects to a greater or lesser extent of those 'economic outcomes', we all have a basic capacity to assess them. Still, you might think that one has to be an economist to assess just how well the economy is working. You need to have expert knowledge to judge whether a given, occasional crisis is simply an unavoidable possibility in the operation of the market economy or whether it might have been avoided, and if so by what kinds of regulations (and how well in general such a regulatory regime might be expected to work). You need to be an economist (so one might think) to recognize that the operation of any institutional arrangement is an issue of long-term optimization, and that the failure to operate 'perfectly' by reference to abstract standards does not show that we are not doing the best that we can. As the economist might insist, 'optimization' involves the best possible response to constraints, and hence any optimum will exhibit some imperfection; and so the mere observation of imperfection is not necessarily grounds for change! In other words, it takes some expertise to distinguish between optimal and non-optimal imperfection—or at least, so the economists might claim. It is not always a straightforward business to assess the quality of putative knowledge, even in the embodied case. Nevertheless, in the embodied case, there is access to additional information—namely, the quality of the products in which that knowledge is

embodied—and in at least some cases, that additional information is relatively easy to assess by the non-expert.

The distinction between embodied and disembodied knowledge, both in itself and in terms of the challenges for quality-assessment, I have treated here as a categorical matter. But of course, the distinctions in both aspects are ones better described in relative terms—as a spectrum rather than an on/off matter. Things that when first discovered appeared as merely ‘interesting results’ turn out to have practical application (and hence an additional dimension for testing their validity). Many pieces of pure mathematics exemplify. By the nature of the beast, we cannot necessarily tell *ex ante* whether the ‘speculations’ of ‘different tribes of philosophers’ will turn out to be useful or not—or whether in the future some application might turn up. Still, it seems possible to locate different ‘tribes’ and different locations of knowledge generating activity along the relevant spectrum. In the remainder of this paper I will focus on knowledge production lying towards the disembodied end—that undertaken in universities (and similar institutions) for the generation of abstract knowledge, mostly without any specific commercial application in mind. There, as I have suggested, the problems of the assessment of quality are rather more complex; and the benefits accruing from more refined specialization, perhaps rather less clear.

5. Disciplines in the Division of Academic Labour

The foregoing is conceived as prologue to a more specific line of enquiry—namely the role of disciplines in the division of academic labour. So it will be useful to summarize what I have said so far in terms of two different interpretations of the work that disciplines do. And I will, for reasons to be explained, focus those different interpretations on the issue of how to understand the relation between disciplines and ‘inter-disciplinary’ work specifically.

One line emphasizes the distinction between production and consumption and sees disciplines as production sites, much as Smith does in the quotation offered at the beginning of *section 4*, above. On this view, the borders between disciplines are sites primarily for communication. Inter-disciplinary work is mainly engaged in translating findings (including both experimental results and lines of reasoning) that are mutually intelligible among practitioners within one discipline into terms that are meaningful within another. Making well-attested findings in discipline A, especially ones that might be used as inputs into knowledge production in discipline B, accessible to practitioners in discipline B is analogous to exchange—and it demands the equivalent of traders/brokers to do this translation work. Because that work involves making disciplinary work accessible to non-experts, the trading activities are complementary with the exercise of extending the relevant knowledge to a general public. But practitioners of nearby disciplines are not just members of a general public: they are a niche audience—one with special interests and agendas of enquiry of their own. So the trader has to have a foot in both camps, talk both languages and understand

the sorts of things that the different audiences will find ‘interesting’.⁸ As in the market for goods, such traders don’t add to productive capacity directly. But they are crucial players in the trading process and Smith in particular would be sympathetic to the possibility that this role would become the ‘*sole occupation of a particular class of citizens*’. That is certainly one interpretation of the work that inter-disciplinarity can do.

The alternative view emphasizes instead the role that disciplines play as institutions responding to the epistemic demands that the division of labour poses specifically in the arena of enquiry. It sees disciplines as mechanisms for certification of ‘product quality’—in this case, for the epistemic validity of the knowledge claims that producers make. That disciplines play such a significant role in the organization of academic labour is on this view simply a testament to the magnitude of the certification problem in the knowledge industry—and specifically in that part of it where the knowledge is not embodied in machines whose working properties can be assessed by the quality of the final products they produce. In this case, *inter-disciplinary* activity needs some independent explanation.

These different views are not mutually exclusive. Indeed, in some ways they are complementary. But they require a bit more development before we can see how.

It is time perhaps for an existential aside.

I have spent the last five or so years of my life participating in and presiding over a PPE (Philosophy/Politics/Economics) program at Duke/UNC. The twin University aspects of this program are not uninteresting, but beyond the obvious point that they reflect a ludicrous ambition on my own part to juggle as many balls as possible, they are not directly relevant here.⁹ This PPE program—like others that have sprung up over the last decade or so, but specifically *unlike* the Oxford variant—has an explicitly integrative and synthetic character, at least in aspiration. The idea is that these three disciplines have an interest-

⁸ Many things that are ‘philosophically interesting’—that is, that philosophers find ‘interesting’—will not be of much interest to economists—still less to political scientists.

⁹ They are though indirectly relevant. If the point about economies of scale in the production of knowledge is broadly correct, one might think that there would be many instances where the banding together of individuals across institutions that are geographically proximate might promise significant returns. If all the economists in the Boston/Cambridge area combined to create the Boston Economics Conglomerate, there would be more experts in more areas of Economics (and the creation of more refined areas) than any single university could offer. This conglomerate would, one might think, be more attractive to graduate students and academic visitors than even Harvard and MIT on their own currently are. And yet, such cross-institutional collaboration seems unusual. And if it occurs it is somewhat covert. Universities are the main institutional entities, and the main reputation bearers—and those institutions rather zealously guard their faculty and their graduate programs as their own. A related issue is why there are not a greater number of more specialized universities (like the LSE say)—why a 14th century institution should have survived the vast increase in the number and correspondingly vastly more specialized ‘men (and women) of speculation’. I note this puzzle. I do not resolve it. But in the Duke/UNC case, it is certainly my strong conviction that PPE as a somewhat specialized enterprise benefits hugely from its juxtaposition of the rational actor political theory strengths of the Duke Political Science department and the moral and political philosophy strengths of the UNC Philosophy department. The whole is more than the sum of the parts!

ing and intellectually significant intersection that deserves to be uncovered and examined—with the rival instincts and dispositions of the disciplines, as well as the overlapping subject-matter, made a matter of deliberate focus. There is a research-agenda aspect and a pedagogical aspect to the juxtaposition but in the context of this paper probably only the former is of interest.

By own background is in Economics; and in some deep sense, I still self-identify as an economist¹⁰ (perhaps in the way that certain children can never quite throw off the afflictions of a Catholic upbringing). But along the way, in a long and rather profligate academic career, I have fallen into dubious company and become interested in unconventional topics and am now probably better described as a PPE-practitioner than in any other terms. Economics, it should be noted, is among the most tribal of disciplines: but because it is defined at least as much by its methods as by its subject-matter, it can be accommodating in terms of what a practitioner seeks to develop an ‘economics’ *of*. The result is that economists have made a practice of occasional imperialist assaults on what traditionally has been neighbouring disciplines’ territory. These imperialist adventures have not, strangely enough, always treated with total enthusiasm by the disciplines being colonized (notwithstanding our impeccable missionary motives). But in any event, the imperialism does not always leave the home discipline entirely unaffected. Economists who stray across disciplinary boundaries come to learn that there are other quite smart people in other disciplines, who do know something—however misguided their methods of analysis might be—and from whom one can, however reluctantly, profitably learn at least something!

The point is that there is an economists’ prejudice—from which I cannot claim to be totally exempt—that what other disciplines offer is territory that can be usefully raided. Being, as we are, in possession of an analytic approach and a set of technical tools in which we have enormous confidence, we can set off on our border escapades with a reasonable expectation of bringing back some intellectually valuable booty. Other disciplines are, first and foremost, an imperialist challenge: when we cross the border in any direction we do so in the conviction that we have something to teach—much less something to learn!

Other disciplines are perhaps less prone to this kind of attitude—and much of what I say here may reflect the peculiarities of the economics case—but I think something of the kind generalizes to other cases. And for more or less predictable reasons. By the time you are a properly credentialed professional available to engage either in imperialist adventures or in defending your own turf, you have spent something upwards of five or six years being trained in the arcane mysteries of your discipline. You have learned to distinguish good work from bad. You have learned to isolate what aspects of an issue are important and what epi-phenomenal. You have—to modify a familiar aphorism—been equipped with a hammer and learned to look for nails. Moreover, your standing in your profession—the reputation you have accumulated and the esteem you derive—depends on your capacity to see nails where others don’t, and/or to hammer them down before anyone else can get to them. And you live in a world

¹⁰ My more mainstream economist colleagues probably have misgivings about my professional credentials—but I have reached the point where this is not a central concern to me.

where such reputation is critical both for its own sake and for the quality of offers (financial and institutional)¹¹ that you secure. Any academic who enjoys even a modicum of success can hope to get invitations to interesting places¹² and enjoy interesting company—all at someone else’s expense; so although we are not perhaps lavishly paid (or at least not as lavishly paid as we like to think we could have been if we had made different career choices)¹³ we live rather wonderful lives, and certainly lives that it is worth expending some effort to construct. So it is perfectly rational to develop our intellectual skills along professionally approved lines—and of course we all spend a lot of time and effort in doing just that.

But of course there is significant ‘path-dependence’ in the development of what is professionally approved. This is a familiar point and is often used to cast doubt on the objectivity and epistemic authority of knowledge. I am disposed to think that that aspect of the story is over-played. The detailed examination of individuals’ capacities at the point of entrance into professions, the scrutiny of work prior to publication in recognized journals, the large professional rewards that accrue to someone who shows that some view widely held within the profession is wrong, all seem to me to work to establish a strong presumption of epistemic warrant. I do not deny (nor do I see how one could) that things within a discipline might go seriously astray, that a discipline could become ‘corrupted’ systemically turning a blind eye to contrary evidence and/or excessively penalizing those who offer heterodox views. We can recognize the logical possibility of ‘Emperor’s-new-clothes’ equilibria in any setting that depends on reputation and esteem to order individual behaviour. I certainly think that professional consensus is often in error—and of course the very idea of progress of knowledge indicates that this is sometimes (perhaps routinely) so. But for all that, I think the odds favour the claim that academic enquiry broadly gets things right, and has mechanisms that encourage error-correction. The fact of professional disagreements in many disciplines in the social sciences—both about method and issues of substance—strike me as evidence less of confusion than of healthy contestation. And though the point often gets lost in the rather fierce battles that we are part of (or observe), it is worth emphasizing the simple analytic point (one I associate with Renford Bambrough) that any genuine disagreement implies much on which the protagonists agree—otherwise, there is mere talking at cross purposes.¹⁴ Hume makes the point that supporters of rivals in political

¹¹ Obviously enough, your esteem is a function of the esteem that your University enjoys and that your specific department enjoys both in the profession and in your university.

¹² Don’t look now. But it is to me striking how often Universities manage to acquire the best buildings in the most luxuriant gardens as work locations and/or conference venues—and I don’t just have Oxbridge and the SCAS Linnaeum in mind.

¹³ And perhaps, just for the record, we ought to acknowledge that we are not, all of us, so very badly paid. A decent economist, in mid-career, might hope to earn \$200,000 a year—perhaps more—which is enough, one might think, to enable you to not have to think much about money (my own personal definition of an adequate income!).

¹⁴ I am reminded in this connection of a story about Joan Robinson introducing a paper highly critical of Friedman’s monetarism when Milton himself was in the audience. She began: “Professor Friedman and I are in vigorous disagreement about the issues at stake in this paper. But we are both economists—and as such we share a very great deal. This is, properly

contests typically exaggerate the differences between the contestants; the same seems to me to be no less true in academic circles.

But back to the division of epistemic labour. Smith's picture of that division is of specialists in more or less amicable agreement as to which issues each will concentrate on, each treating the others' findings as more or less authoritative and the others' motives and capacities with respect. There is little room in Smith's picture for passionate disagreement, contempt for others' work, for obfuscation and lack of communicability—things that we often find in inter-disciplinary engagements (and sometimes in intra-disciplinary ones).

Nor is there any suggestion of the strong centripetal forces that disciplines exert. For the same incentives for professional esteem (respectability) that are essential for establishing the endorsement of any individual's work by (somewhat) independent recognized experts tend also to draw practitioners inward to matters that are recognizably 'mainstream'. At least in disciplines with a strong identity, the remark: 'Oh, X. He's not really an economist (philosopher)' is a remark of dismissal. It is a judgment that most are quietly alert to avoid.

And these centripetal forces have an intellectual aspect—they have the effect of insulating disciplines against the possibility of dependence on other disciplines. So, for example, the tendency in economics to appeal to what economists see as very 'weak' and uncontroversial normative claims in the development of policy (or institutional) recommendations; or the somewhat analogous attempt by moral (and political) philosophers to finesse matters of fact by appeal to 'ideal theory'. In the interface between economics and moral philosophy specifically, it seems clear to me that any substantive normative claims will involve purely normative elements alongside best available accounts of the way the world works. In that sense, economics and ethics need each other. And one might think that, if both parties recognized that fact—if they carried a picture of the overall intellectual enterprise as characterized by the division of labour—they might fashion their efforts so as to make the results of their enquiry maximally hospitable to the necessary input from their disciplinary neighbour. Strangely, the disciplinary representatives are not much prone to saying, in response to some substantive question: 'Well, we can't say much about this without major input from the economists or the philosophers or the political scientists.' Well-acclimated (one might add, *profitably* acculturated) disciplinary prejudices incline practitioners to think that their own discipline has far and away the most to contribute—that 'X is *really* an economic (or philosophic or political) issue'. This is a kind of dialogue where 'or' dominates 'and'.

It goes without saying, I think, that the work of brokering profitable exchanges across disciplines requires a person to have a foot in both camps—to be able to understand from a more or less internal point of view what both the economists and philosophers are saying. You cannot translate from Spanish to English on the basis of a mastery of Esperanto! But here there is more at stake than just translation. An economist who is philosophically savvy, who is

construed, a family argument." (Or something of the sort. The occasion was reported to me, long after the event, by a Cambridge undergraduate of the time who was also present.) The sentiments seem right to me.

alert to conceptual issues in her discipline, can see things that are both economically and philosophically interesting. A philosopher who is alert to feasibility constraints—whose instinct is for example to see the brain as an organ with an essentially limited capacity, or political institutions as mechanisms for coordinating the activities of intrinsically morally imperfect agents—will tend to foreground issues that might otherwise be overlooked. It is in that sense, I think, that the inter-disciplinary translator becomes more than a translator: she becomes a substantive player in the various fields across which she moves. That is the sense in which PPE is more than the sum of its parts. Or at least, so I hope!

6. Conclusions

The aim of this paper has been to mobilize Adam Smith's treatment of the division of labour in relation to the production, consumption and exchange of knowledge.

That mobilization has two aspects. One deals with the epistemic demands that exchange makes on participants in the trading nexus. The advantages (manifest and very substantial) that are associated with specialization and the division of labour depend, I have argued, not just on the 'size of the market' but also on the capacity for market agents to solve the principal-agent problems associated with information about product quality.

The other deals with increasing returns in the provision of knowledge itself—treating knowledge creation (as Smith does) as just another example of specialization and exchange.

These two aspects come together, however, in relation to the epistemic demands associated with assessing knowledge quality. These demands differ somewhat according to whether the knowledge is embodied in products (and machines) or whether the knowledge is an object of desire for its own sake. In the latter case, the test for the epistemic authority of the 'knowledge' often lies in the testimony of 'experts' (though in some case where the capacity for 'recognition' is wide-spread, little expertise may be required.)

My claim is that in the academic arena disciplines play a critical role as institutions for meeting the epistemic demands that the division of labour creates in the 'knowledge' case. In that sense, I believe disciplines are indispensable. But any division of labour presumes a distinction between producing and consuming, and a role for *exchange*. And inter-disciplinary communication is part of that exchange activity.

I have though suggested that disciplines come at a cost—that of creating centripetal forces that divide them both institutionally and intellectually. So the exercise of inter-disciplinary communicability becomes itself a niche in the intellectual division of labour, one requiring important specialist skills.

But the role of the translator/trader here is not just one of aiding communicability. As she imposes on discipline A the intellectual vision of discipline B—and vice versa—the translator becomes a potential producer. There are lots of quite interesting pickings both along the borders between disciplines and in

the distinctive perspective that comes from the capacity to eye the world simultaneously through different lenses.

It should be clear that I do not regard disciplinary specialization as a blot on the academic landscape. Many who extol ‘inter-disciplinary work’ seem to. I think that those who oppose ‘excessive specialization’ are often insensitive to the huge advantages that specialization brings. But there is a highly desirable niche market for the genuine inter-disciplinarian; and it may be that that market is larger than current academic practice would suggest.

Bibliography

- Baumal, W. (1967), Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis, in: *American Economic Review* 57(3), 415–426
- Prebisch, R. (1950), *The Economic Development of Latin America and its Principal Problems*, New York
- Smith, A. (1981[1776]), *An Inquiry into the Nature and Causes of the Wealth of Nations*, Indianapolis
- Wilson, B. (2008), *Swindled: From Poison Sweets to Counterfeit Coffee*, Princeton