

7. Integrity

We've been using the word *integrity* as a blanket term for all those qualities that give epistemic value – value as worthwhile knowledge, productive of further knowledge – to a conversation process. Whether we think of a single experiment or a whole scientific discourse, a case at law or a whole legal system, a dialogue between friendly adversaries or a bitter controversy in the newspapers, it's the *integrity* of the process that gives knowledge-value to the argument that develops. In the context of argument and conversation, that is what we mean by integrity in this special sense, and we choose that word for its consistency with ordinary usage. But now we ask, of what does such integrity consist, and how is it to be judged?

This is an urgent and practical question since information – suggestions to believe and act in some fashion – bombard us constantly from a variety of sources, none wholly trustworthy and many intentionally deceptive. All sources beyond personal experience – and even that perhaps – may be considered arguments of a sort: outcomes of a process whereby competing suggestions are made to engage each other, and arrange themselves, into some cognitive pattern or structure. Everything we think we know, all the information we entrust ourselves to, may be seen as the product of a conversation, and will be trustworthy just to the extent this conversation has what we call integrity. But to name a thing does not explain it, of course. We must still consider just what it is that makes a conversation worthy of our trust as a source of knowledge. What factors or properties contribute to a conversation's epistemic value?

To begin with, it's clear that integrity must be a procedural notion – a judgment on the competence and trustworthiness of some informant's internal *process*. We don't know what *truth* is, but have a pretty fair sense of what it means to be ignorant, or to tell lies. Having no direct access to the state of things in themselves, our knowledge of them can only be the outcome of some perception or investigation; and our confidence in that outcome can only rest on our evaluation of that investigation. Of course, we can compare the findings of one investigation against another that we trust more. We can ask for second opinions, as it were. But finally, it always comes down to a question of which process we trust the most.

For example, in a criminal trial, the law bases its judgment of an accused's guilt or innocence on the outcome of a "fair" trial. There is no way to decide whether a verdict is correct, except by evaluating the procedural integrity of the trial itself – selection of the jury, value of evidence presented,

competence of the advocates, propriety of the judge's rulings, etc – or else by appeal to some external process that we trust more – e.g. retrial in an appellate court, or a subsequent confession with details that only the guilty party could have known. Sooner or later, we will either have to abandon reason and go with the most congenial conclusion, or else rely on the process whose “integrity” we have the most reason to credit – or the least to discredit.

In science too, the experimental test of an hypothesis is not conclusive or even relevant without a similar judgment of procedural soundness. We have to decide whether the experiment tested what it was supposed to test, gathered complete and accurate data, analysed these correctly, interpreted the results judiciously, and so forth. Even the “hardest” science rests, ultimately, on a foundation of critical appraisal of integrity in experiment and argument. Indeed, “hard” science is hard, precisely because the integrity of its procedures are relatively easy to verify. But, in the last resort, even hard “facts” are only considered such with reference to the integrity of the observers who bear witness and vouch for them. My confidence that the object on my desk is a mug and not a raven hangs on a judgment about the reliability of my eyesight.

7.1 Pathologies of Conversation

How, then, do we evaluate the integrity of a conversation? It's easiest to come at this question from behind by enumerating all the well-known pathologies that sap our trust in a conversation's results. Evaluation of a national intelligence service makes a clear example of our procedure here. The paramount mission of such an organization is to prepare reliable briefings on the world's state for the chief executive and his privy council (by whatever name it is called). The apex of pragmatic conversation is in this executive's hands: He has the final decision on what will be done. By contrast, his intelligence service has an epistemic mandate: to inform executive judgment by submitting valid, useful reports and briefings as a basis for decision.

Nearly always, however, the raw information in these reports will need to be interpreted before it becomes useful. It is one thing to count items of hardware in some other country's arsenal; it's entirely another thing to assess that country's intentions and military capabilities relative to one's own in a specific situation. And for that matter, the sometimes cryptic pictures from cameras in the sky, and the sometimes fudged, prevaricating or conflicting yield from field agents, will also need to be interpreted. The ultimate client for intelligence material thus faces an immense Nietzschean predicament: The executive must take his decisions based on understandings that owe everything to interpretation – hopefully guided by some degree of measure, harmony and reason. In this book's language, we say the value of the intelligence provided him will depend (by definition) on the integrity of conversation behind it.

Specifically, the reliability and usefulness of his intelligence reports will depend on qualities of the intelligence processes – of information gathering, collating, evaluating, interpreting, summarizing, presentation, etc. – and on the distortions to which these processes are vulnerable. The simplest way to explain what is meant by integrity of conversation is just to enumerate the factors that make any intelligence organization (including one’s own brain) work well or badly.

Witness

A first area of conversational competence might be called *coverage* or *witness*. To have significant epistemic value, a conversation must in some sense bear witness to, or “cover” the matter of which it speaks, in much the same sense that a witness testifies, or that journalists cover a story. The intelligence service could scarcely prepare useful briefings without on-site sources; and those briefings would be worse than useless if those sources were double agents. In general, to enumerate how witness can fail, we need only consider how a critic might seek to discredit it on some given issue. Most obviously, a critic might point out that the sources had no contact, or very limited contact, or only duplicitous contact with the matter in question; that they were not in a position to know very much, or were unlikely to reveal what they knew.

This is why the inevitable self-interest of persons best positioned to have first-hand knowledge of any matter has such serious consequences for the epistemic value of our public conversation. Specialists bear witness about their specialties, but do so always from a viewpoint of personal and professional interest. This being so, public knowledge cannot mean what was once hoped. What it can mean will be discussed in Chapter 9.

Interpretation

As we have seen, all understanding is *understanding-as*: an assimilation of what is present to what is remembered, construed or imagined. A conversation may have good sources in some area, yet lack the skills or resources to connect its present perceptions to previously acquired concepts or experience. In other words, it may fall down on interpretation.

Continuing with our example: Certainly, a head of state should doubt the value of his briefings, if he suspects that his intelligence chiefs are telling him what they think he wants to hear. He would be annoyed if the reports were presented to him as undigested detail, rightly demanding to be told a relevant story, not pestered with meaningless facts. Thus, two further pathologies of intelligence conversation are wishful thinking, and interpretive timidity – the inability or refusal to assume the risks of putting a story together.

The rub, of course, is that the interpretations available (alternative hypotheses, as we may think of them) carry different risks associated with their acceptance or rejection. As in statistics, the risk of accepting a poor interpretation must be compared with the risk of rejecting a good one. These risks must figure into the structure of argument amongst competing understandings, on which the actual decisions are based.

Argument

Here we get ahead of ourselves, since the nature of polyphonic argument will be discussed in the next chapter. We have chosen to discuss integrity first because, without this concept, the notion of reasoned argument cannot have meaning. But there is unavoidable circularity here, because the integrity of a conversation depends in part on the quality of its argument. Certainly, part of what we mean by integrity in a conversation is that its arguments be conducted respectfully and honestly. On the other hand, it is not possible to say what is meant by integrity of an argument without reference to that of the conversation as a whole. The reason is that good argument is more than a conflict of opinions. It is *witness*, since we are brought to perceive the matter from our opponent's perspective. It is *interpretation*, since we are continually interpreting our opponent's understanding in light of our own, and re-interpreting the issue in the light of arguments advanced. It is *commitment* (see below), since our claims to seriousness or veracity represent small commitments that we make to our interlocutors. In fact, we will have to say that argument permeates every phase of conversation, and is itself permeated by all the others.

Nonetheless, thinking of argument in its narrow sense, as the adversarial aspect of conversation, we can give a correspondingly narrow account of its integrity: Argument of integrity is sharp, but open and courteous. It need not be friendly, but cannot afford to be contemptuous of any viewpoints taken seriously at all. Above all, it takes care to include and deal thoughtfully with every viewpoint that seems too important to ignore. It can be seen as a social game, like bridge or golf. As such it manages to be adversarial and collaborative at the same time: It is adversarial, naturally, in that each side hopes to see its opinions prevail. Yet it remains collaborative also, in that victory will be meaningless if the conversation is not left sufficiently intact to sustain a common commitment to its own outcome.

Commitment

After the range of opinion and argument on some matter has been worked up and then allowed to settle down into a stable structure, a final task remains. The intelligence exercise is not complete until its briefing documents and

presentations are prepared and delivered. Thus, even with the argument spun out and documented, there remains a phase of commitment: the organization's actual statement to its clients.

Just here, perhaps, is the greatest weakness of the entire process: Typically, it will be the Intelligence Director (let's call him Q) who accepts the end product of his organization's efforts, presents it to his clients, responds to their questions and endures their scepticism. Typically, it is Q who controls the other phases of the process. Thus, it is all too likely that interpretations and even hard facts that Q would feel uncomfortable presenting will be excluded from the outset – so that Q's clients, and even Q himself, are told nothing that would conflict with what they already expect to hear. In this way, the purpose of the whole exercise can be defeated by a failure of intellectual integrity at the top. This may happen in various ways: There may be a single dominating obsession that prevents a balanced regard for other concerns and issues. There may be pathological fear of closure, or of the responsibility it entails. There may be pathological haste to put all doubts to rest. There may be simple laziness or bureaucratic self-protection: the wish to avoid any conclusion that would require action. Or, on the contrary, there may be desire to justify action of any kind.

* * * * *

What then should Q's boss demand? He should not ask Q to guess the future, but to convey the structure of an argument of integrity on what the future is likely to hold. As suggested in Chapter 4, that structure will be a better basis for decision than a "best guess" could be – even a well-informed guess. The reason is that a briefing on the structure of argument prompts the decision-maker to consider scenarios other than the most likely (or the one Q believes he wants to hear), and to hedge his bets accordingly.

7.2 Conditions For Knowledge

We can look at the question of conversational integrity from the other end as well. We can ask, under what conditions will diversity of competing understandings amount to useful public knowledge rather than confusion and strife?

Ideally, for their discussion to amount to public knowledge, the interlocutors should at least come to understand each other – understand where their adversaries are coming from. They may not come to any agreement; they may remain opponents or deadly enemies, even; but if they can at least grasp each other's claims, and recognize the basic humanity of those who make them,

significant knowledge will have been gained. At the minimum, a contribution to fruitful negotiation will have been made when each learns to accept the others' reality as an aspect of his own. War creates this acceptance eventually, when both sides have bled sufficiently. That is why war remains the arbiter of last resort in international affairs. But it is better when adversaries can recognize each other's rights to existence and cognitive sovereignty from the outset.

The crucial condition for fruitful public knowledge must be that the disputing parties feel some stake in the quality of their debate, rather than merely in its outcome. So long as they can do this, rival laboratories working from different paradigms will advance science. Rival advocates arguing opposite sides before a judge and jury advance the cause of law. Intellectuals of various stripes will gather about them a clerisy of well-informed and active citizens. Without mutual concern for the integrity of the process, useful knowledge may still accrue – but only to cooler third parties, perhaps of the next generation, who can get their heads around the quarrel and proceed from there.

In effect, there are at least four reasons why the structure of an argument of substantial integrity may be recognized as a kind of truth:

- (1) First, as has long been understood, a review of the free, well-conducted argument around any matter affords the best possible opportunity for independent judgment as to the relative merits of each position. In particular, it affords opportunity for judging the intellectual integrity of each position on its own. You may prefer one position to the others because it is most convenient to your aims and wishes, or because it seems best to satisfy the available evidence. That is entirely up to you. But the argument as a whole may be considered public knowledge because anyone who follows and gets his head around it will have substantial grounds either to rest in uncertainty or to espouse one position as his own.
- (2) Second, the idea of “polyphonic” truth – truth as a structure of reasoned argument – reduces to flat, propositional truth in the special case that argument collapses. When no one wishes to dispute a given interpretation, or has yet found an effective way of doing so, that interpretation is, to all intents and purposes, “true” in the classical sense. Thus, the notion of “truth” as a structure of argument is fully consistent with the existence of “fact,” understood to mean “what is jolly well the case.” To speak of “fact,” what is necessary is not just overwhelming consensus, but a judgment of fatuity – absence of serious experience or argument – in whatever dissent exists. Thus, *pace* some of Nietzsche's

more extreme followers, we can still take it as “fact” that the moon is not made of green cheese. Biologists, though not theologians, can take Darwinian evolution as “fact,” because the scientific case against it is fatuous.

- (3) Polyphonic truth serves as well, and perhaps better than classical truth as a point of departure for further study and argument, and for the accumulation of personal and public knowledge. Differing understandings of any seriousness have their roots in different experiences, or in different approaches to what would otherwise be fairly similar experiences. Such understandings or approaches may or may not have anything useful to learn from each other, but they at least offer a range of possible starting points to the uncommitted. There is, at worst, an embarrassment of interpretive riches. This need not be a problem, if you feel entitled to start where you happen to be standing, and to follow your nose where it leads.
- (4) Amongst disputing parties or factions, the results of an argument of mutually-recognized integrity provide the best available basis for negotiation, politics, and jointly acceptable settlement or decision – whether or not the argument converges. Even when the argument does not converge, its structure can be discerned and accepted as a common reality from which politics and negotiation must proceed.

We should not be surprised that polyphonic truth avails in all these ways, because it’s what we have always had to make do with when flat truths could not be had.

7.3 A Definition of Reason

We can now say more precisely what *reason* means: **To be "reasonable" about some matter is precisely to be accountable through conversation of integrity to the perceptions and concerns that bear upon it.** Reason itself then is a climate and practice of being reasonable.

This definition is in good accord with ordinary usage. For example, your boss can demand to know why you handled a certain case in a particular way, ignoring available expertise, established policy, current thinking, or known stakeholder concerns. A reporter or a Member of the Opposition can ask a Minister the same questions. A hostile lawyer in a malpractice suit can ask these questions of a doctor. In every case, the incumbent will suffer embarrassment if he cannot show that he had considered the relevant issues and arguments, and had good reasons for believing or acting as he did.

Beliefs and choices remain accountable to the relevant conversation of integrity on a matter because, or to the extent that, its structure of argument cannot be safely ignored. Taken as a whole, the structure of argument constitutes a public understanding, however conflicted. Ignorance or negligence in the face of this understanding will lead to public embarrassment. When an issue lands on your plate, neither law nor custom, and nature least of all, will absolve your mistakes on grounds of ignorance.

But the crucial point is that the concept of accountability to relevant conversation involves a presumption of that conversation's basic integrity. Thus, it is not merely that people should respect each other's concerns, or that they should bear honest witness, etc. but that they **must** do so with some reliability or see their affairs bog down in confusion and strife. Norms like these are "laws" of the logic of conversation because their trespass leads to nonsense, just as happens when conventional logic is trespassed. The latter cannot compel people to reason correctly. It merely shows how contradiction follows when they fail to do so. In quite a similar way, conversation's logic does not enforce an ethic of discourse. It merely tells us what to expect when violations of that ethic come to be expected. The conditions of post-modern society amply confirm its central prediction that failures of conversational integrity accumulate in vicious circles: The more such failures there are, the more there will be. As this happens, our conversations keep getting louder, more hostile and more volatile, while their knowledge value falls toward zero.

Once the integrity of a conversation is hopelessly compromised, it can no longer serve as a reliable standard of judgment, and its participants find themselves accountable only to superior power. Without some standard of integrity that people can be expected to share and meet, there is no possibility of holding anyone to account for a lapse of honesty or judgment. Conversely, in holding a person accountable for his beliefs and choices, the integrity of a relevant conversation is asserted and defended. To the extent this becomes impossible, reason in public affairs is at an end.

It's like setting a clock. One clock can be set to another that is known to be accurate, or it can be set to the average of a few clocks that were once synchronized, and are known to gain or lose by some random amount. But, given n clocks whose settings are entirely haphazard, the setting of the $n + 1^{st}$ clock must also be arbitrary. If no information about the correct time can be read from the clocks in one's possession, then any guess is as good as another.

As a corollary, it will be clear that an epistemology of conversational integrity is not less demanding than the classical epistemology of "objectivity", but much more so. Classically, the epistemological obligation was simple: One could claim to hold a judicious opinion through a sincere effort to be right. By contrast, in conversational epistemology, being "right" is not nearly good

7. Integrity

enough. Now the obligation is to grasp the situation “in the round,” so to speak – from the viewpoints of the other players as well as your own. Such multi-faceted understanding requires formidable efforts of detachment, imagination and self-discipline; and proves much more difficult in practice than the objectivist requirement “to get the facts.”

Finally, it will be seen that conversational integrity is a form of “situation ethic,” undercutting the authority of law and principle. Deriving policy from principle is hopeless, because principles are ideals and therefore permanently in conflict. They cannot be taken as absolute commandments or prohibitions, but merely as competing poles of suggestion. The role of principle is not to determine action, but to mark the extremes of possibility and value. In the event, people are (and should be) actuated much more by particulars than by generalities. The abortion issue is a perfect example: On the theoretical plane what is involved is a choice between irreconcilable values and understandings. But the decision whether to abort an actual pregnancy is a choice between two cruelties, only recognized as such through a contention of understandings and moral judgments – with some concrete decision, eventually, based on particulars of the case.

Principles create the field of tension in which specific choices are made; and they are not to be held, so much as wrestled with. Without the cognitive tension (tensegrity) that principles afford, there will be conflicting interests, but there can be no genuine concerns regarding ethics or values. There is room for horse-trading, but not for intellectual or moral deliberation. The power dimension of choice remains, but its cognitive dimension is lost. The policies and plans formulated in such a climate of moral bankruptcy will probably be crude and self-defeating, because they will not have considered the complexity of the situation and the range of requirements to be addressed.

Thus, we cannot expect to deduce concrete choices of any complexity from *a priori* beliefs and values. These are usually in conflict, and therefore incapable of giving a coherent answer on any concrete matter. Rather, we use beliefs and values to inform our decision of the concerns at stake. Given the uncertainties of life, and our differences of viewpoint, we can never say whether a particular choice was the best possible in any universal sense. Until all the results are in (which may take forever) we cannot even begin this argument. Thus, we will never be able to agree how good a choice was, but should at least be able to judge the quality of the conversation in which it was framed and made. We can usually tell how *sound* a decision was – how responsibly it was taken.

To anticipate the next chapter somewhat, we find the word *argument* used with three distinct meanings:

- (1) It is the name for a kind of discursive process that seeks to alter cognitive commitments whether by reasons or by rhetoric.
- (2) It may refer to the tensegral, structural result of such a process, as when we speak of “getting our heads around an argument” (as a whole).
- (3) But finally, it is also the word for representations put forward in the course of an argument process to get people to “change their minds” – to perceive or feel differently about something.

It is the last meaning that concerns us now, and our question is this: Granted that no reasoned argument can *force* an individual to change his mind, when does it *deserve* to do so? Attempts at cognitive influence are often made in ways that have nothing to do with reason; but reasoned argument seeks to change your mind by cogently explaining why you *ought* to credit a certain cognitive commitment or viewpoint. Clearly without some ground for normative judgments along these lines, the concepts of reasoned argument, and of reason itself, become so weak as to be meaningless.

Now, classically, the ultimate source of reason’s authority has been the concept of absolute Truth¹. For who would want to believe what is False? It might be expedient to get someone else to believe it; or one might fall into self-deception from emotional necessity. But the undesirability of such lapses appears self-evident: It seems obvious, the ancient philosophers thought, that people ought not to play such tricks on others, and certainly not on themselves – that everyone ought to manage his cognitive life so as to commit himself, and suggest to others only what is True. Classically, we ought to “follow the argument wherever it goes,” and believe the argument’s valid conclusions.

Our difficulty is that in polyphonic domains, where the concept of absolute Truth becomes untenable, the normative force of argument must find a different basis. We cannot derive an “*ought*” of cognitive commitment (even our own, let alone other people’s) from an over-arching, Socratic commitment to absolute Truth in domains where no such Truth exists. However, the cause of reason is not lost if the required *ought* can be derived from some other source. It may still be possible to refer the normativity of well-reasoned argument to internal qualities of the conversation process itself. Indeed, some

¹ The concept of formal validity does not affect this point. In formal domains, the validity of an argument is identified with the truth-value of the following proposition: “The conclusion of this argument has been derived correctly from its premises, under the pre-established transformation rules.” Thus, the normativity even of mathematical argument still depends on a truth-judgment – in this case, about an assertion of formal correctness.

version of “integrity,” or “discourse ethics,” seems to be our only hope for worthwhile notions of reason and public knowledge, beyond the conventions of this or that particular knowledge-culture.

If this is so, it will be apparent that corrupting the integrity of the public conversation must be a very dangerous form of pollution, since it erodes the possibility of public understanding, and of reasoned public choice, and thus makes reasonable government impossible. Since reasonable government is a necessity of tranquil social existence, the normativity of reason seems to follow: Unless we welcome chronic violence, we can say with confidence that we *ought not* to corrupt our conversations, nor allow others to corrupt them. We can see that public leaders who do so are undermining the basis of their own authority. They may cling to power, and will probably do so for some time to come; but they will find their options narrowing, the tides of violence rising, and their positions increasingly precarious.

Categorically, then: *The only valid reason why one opinion may be considered superior to another is that it derives from a conversation of superior integrity.* Without such judgments of relative integrity, one opinion is always just as good, but never better, than any other. It is only with reference to an idea of conversational integrity – comprised of deductive rigour, scrupulous experimentation, intellectual honesty, critical judgment, or all of these together – that we can explain:

- why some interpretations are stronger than others, and more deserving of epistemic confidence;
- what it means to offer reasoned arguments (as distinct from merely persuasive ones) for a particular viewpoint;
- what we could mean by *public* knowledge, as distinct from the *local* knowledge of some particular culture;
- what we could mean by a rational public choice, taken in the *public* interest, as distinct from one serving some governing coalition that has succeeded in getting its own interests identified with those of the community as a whole.

We can no longer speak of public truth or knowledge in anything like the classical sense, but only as a tensegrity of argument. However, through the notion of integrity, we can at least explain why we place our cognitive bets the way we do, and hold commitments to one position over another. We can try to be fastidious in pinning our cognitive preferences amongst alternative understandings to judgments of their integrity. This is what we can mean by

reason today; and I suspect it is what people have always meant in practice.