Testimonial Beliefs and Epistemic Competence[1]

David Henderson

Forthcoming in *Nous*

1. Introduction

Much that one believes is dependent, wholly or partially, directly or indirectly, on testimony—on accepting written and spoken reports. Many such beliefs are arrived at in epistemically good ways, and are objectively warranted. But, in virtue of what are they so? I focus here on an issue regarding which there seems to be some disagreement: whether one may be, or typically would be, warranted or objectively justified in accepting testimony without monitoring the trustworthiness of one's interlocutor.[2] Some writers, for example, Fricker (1987, 1994, 1995) insist that testimonial belief formation is only warranted when the agent forming beliefs monitors the interlocutor for trustworthiness. Fricker argues that those who would not demand such monitoring wrongly license simple gullibility. It is not hard to find authors who seem hesitant when it comes to requiring that one "monitors" or examines, or gauges the likely trustworthiness of interlocutors. Yet, not surprisingly, it is *not* easy to find accounts of testimony that would clearly license gullibility. My goal here is to articulate what *ought to be at issue* in demands for monitoring. I will argue that some form of monitoring *is* requisite to warrant testimonial belief—and will endeavor to say some informative things about the needed monitoring and how it is routinely and tractably managed. Those who resist the demand for monitoring typically hold that one has a defeasible default entitlement to accept testimony, unless one has reasons not to do so. Once one gets clear on what should be demanded by way of monitoring, and on what should be intended by the suggestion that one has a defeasible entitlement to rely on testimony, there may be considerably less disagreement then it might seem on the face of the present debate. If it should turn out that all parties to the contemporary

discussion agree with the position I develop here, I would be happy simply to have clarified a rare philosophical position on which others uniformly agree. At the very least, there are philosophers who would be surprised should there be such agreement.

I should acknowledge at the onset that what I have to say about each of the competing approaches may strike some as somewhat disorienting, perhaps even as distorting. I believe that any "strangeness" to my characterization of the debate arises from my insistence that the competing positions be read as proposing an understanding of what cognitive processes make for warranted or "objectively justified belief." *Being* objectively justified (or warranted) in holding some belief requires more than merely *having* justification for that belief. It requires the right kinds of cognitive processes also be in play upon the justification one has—and so I take the debate to concern what constitutes the epistemically appropriate kind of cognitive processes. I argue that those who suggest that we need not monitor for trustworthiness should yet be understood as supposing background cognitive processes with much the same effect as the justificatory processes envisioned by partisans of monitoring. At the same time, partisans of monitoring allow monitoring to be accommodated in the cognitive background, in ways that should be congenial to the "opposing" camp.

Because my process-centered understanding of objective epistemic justification, or warrant, informs my approach to the issues of concern in this paper, I ask the reader to indulge me in a short exposition of how I conceive of objective justification. I think of this as a matter of a matter of exercising a kind of *epistemic competence* or epistemic virtue, which, for normal adult humans, would be captured in a descriptively constrained ideal model of adult human belief fixation.

2. What Makes for Epistemic Competence

The appropriateness of processes—what processes make for objectively justificatory belief-fixation—cannot be understood apart from the kind of agent whose processes are under consideration. The claim that one epistemically ought to fix beliefs using certain processes implies that one can do so—that *agents with the characteristic endowments of the relevant class of cognitive systems can* implement

the processes in question. This is an important insight running through the work of Cherniak (1986), Harmon (1986) and much naturalized epistemology.

The notion of "can" that constrains standards for objective appropriateness of processes is different from that in play when one judges that, given some *particular* agent's poor training and resulting dispositions, that agent could not but think along certain lines. If an agent's training, motivation, or the like, are "substandard," that likely does limit what that individual realistically can manage—but typically that will not excuse the agent or lessen the epistemic standards one is inclined to apply. Rather, one's epistemic standards for objectively appropriate cognition, one's everyday epistemic evaluations, are conditioned by understandings of *what agents of the relevant sort can manage with good training and reasonable motivation*. The default class of relevant agents is normal adult human beings. With training, and in conditions free from special duress, distraction, fatigue, and the like, humans can come to manage or implement certain processes, and will prove systematically incapable of managing others. Fitting epistemological standards select from those dispositions that humans *can* come to possess (with congenial training) a set of dispositions to processes that would jointly significantly further the satisfaction of our epistemic end: the production of true belief systems, of understanding. Of course, any actual human being will have other dispositions as well—dispositions that degrade one's epistemic performance, inducing divergence from competence (dispositions to fatigue, distraction, and the like). These are dispositions to various sorts of epistemic performance errors.

I have just now moved freely from talk of forming beliefs in ways that make for objectively justified belief, or warrant, to talking of processes that constitute human epistemic competence. What is the nature of the connection here? The central idea is that our concept of object epistemic justification centers on the idea of agents—typically normal adult agents—employing epistemically optimal or satisfactory sets of cognitive processes to form beliefs (that are thereby objectively justified). "Epistemic competence" is the term I employ to characterize the set of such yoked processes. At any given point, our understanding of epistemic competence, our understanding of what processes of belief formation make for objectively warranted or justified belief, is an idealized standard by which we gauge the epistemic

quality of our own and other's belief production and maintenance. (Instead of "epistemic competence," I might have written of "epistemic virtues" (Henderson and Horgan, forthcoming 2).)

An understanding of normal adult epistemic competence then provides epistemic standards fitting to humans. This understanding (these standards) is (are) both descriptively constrained and ideal. It is descriptively constrained because what is fitting is limited by what cognitive processes and dispositions can be acquired by humans. At the same time, what makes for objective justification is ideal—our standards reflect our best understanding of the *select* set of epistemically conducive processes and dispositions that humans can acquire (Henderson 1994). Such standards are open to revision as we come to refined understandings of the processes that humans can acquire and their tendencies. Conforming to extant standards does not, of itself, make for objective epistemic justification—as there is a difference between being objectively justified and merely conforming to one's present standards for gauging objective epistemic justification. Such is the force of talk of 'objective' versus 'subjective' justification.[3]

2.1 Reliabilities

Reliability is certainly a very general feature of processes that relevantly conduce to the satisfaction of the central epistemic end (the production and maintenance of systems of true beliefs). Here I will discuss two forms of reliability, *global* and *transglobal* reliability. For my own part, I have come to be convinced that the less familiar form (transglobal reliability) is the most significant when thinking of objective epistemic justification.[4] But in this paper I discuss how monitoring contributes to both. I will also discuss the character of the monitoring or regulating processes so as to make the case for their tractability for normal adult humans—indeed they are typically managed in the background, in a way that is relatively cognitively costless. This will add up to a good case for the view that processes that monitor for trustworthiness constitute a component of the set of processes making for normal adult epistemic competence.

Reliability is always relative to an environment. A process might be reliable in that local environment through which an agent passes at a time—such a process then has *local reliability.* An

agent's global environment includes the set of local environments that the agent might readily confront. A process might be locally reliable without being globally reliable, or might be globally reliable without being locally reliable. The reliability of processes with respect to the agent's global environment—call this *global reliability*—is clearly a feature that contributes importantly to one's epistemic success. Accordingly, global reliability has come to figure prominently in discussions of what processes give rise to objectively justified beliefs.[5] I later discuss how processes that are sensitive to indicators of degrees of trustworthiness, and that regulate the acceptance of testimony accordingly, contribute to the global reliability of one's belief fixing processes.

However, a striking, if extreme, sort of scenario indicates that global reliability is not strictly necessary for the sort of epistemic competence that makes for objective justification. Imagine an agent being subject to a powerful evil demon. No process would be reliable in such an environment. But most find that they are strongly inclined to judge that the notion of objective appropriateness and justification has application here—that an agent beset by such an evil demon could yet be objectively justified in believing. Sosa (1991) has termed this the "new evil demon problem" for reliabilism.

This points to the second epistemologically valuable form of reliability. Transglobal reliability is had by processes that would be reliable in a wide range of epistemically relevant possible global environments—where, an environment is epistemically possible if it is compatible with the agent having experiences of roughly the character of those we have (Henderson and Horgan 2001 and forthcoming). The epistemic value of transglobally reliable processes is easy enough to understand. It is a fact of epistemic life that one must go about one's epistemic chores, one must fix beliefs, even as one remains rightly uncertain as to the exact character of the world in which one labors. Intuitively, compatible with experiences of the rough sort one has, one could inhabit a range of possible environments, and one is best off, one is *epistemically relatively safe,* to the extent that one employs processes that would be valuable in a significant range of such environments. Think again of an agent in the demon infested environment. Suppose that that agent employs processes that are pretty much exemplary of what we would like—with all the checks and safeguards that we expect when engaging in careful observation, inductive reasoning,

and so on. Of course, this will do that agent no good in the pursuit of the hallmark epistemic ends. For, in that inhospitable environment, those—indeed any—processes would fail. Still, insofar as those processes are of a sort that would yield largely success in many epistemically possible environments, one is inclined to judge that the failures are to be charged to the perverse environment, and not to the agent. Fine agent, fine processes, just extremely lousy and unfortunate environment.

Both global and transglobal reliability are matters of degree, and gains in global and transglobal reliability commonly come at an epistemic cost. After some reasonable threshold, gains in robustness or reliability may be offset by costs in terms of two other familiar epistemic virtues of processes: power and speed. In section 4.3, I argue that the monitoring of interlocutors for trustworthiness brings obvious gains in both global and transglobal reliability. Further, I suggest that these gains can be gotten through tractable processes that do not incur high costs in terms of power and speed. In section 4.4, I indicate how such monitoring might well be managed rather automatically. All this adds up to a strong case for an epistemology of testimony that enshrines some form of monitoring.

2.2 Components and Not Missing a Trick

Epistemic competence (and thus objective epistemic justification) is a package deal—constituted by a package of complementary processes that jointly possess the desirable balance of reliability, power, and speed.[6] A process can be a component of the package or set constituting epistemic competence and yet, without the use of complementing processes that refine its working in one of various ways, that process may not be by itself epistemically desirable, or not satisfactorily so. The results of one component of epistemic competence in isolation, without the benefit of refinements provided by other components, may count as a serious performance error. Consider the processes that jointly make for competent generalization from samples.

First, consider a simple process that generalizes from the observed prevalence of some property in a sample. Certainly such inductive inference from samples has some place in the package of processes that makes for human epistemic competence. But, as epistemically useful as generalizing from samples

can be, it is useful in the attainment of the central epistemic goals only to the extent that it is combined with processes that inhibit the generalization from unrepresentative samples. That is, the simple process of generalizing from samples must be yoked with tractable processes that monitor prospective samples for bias and inhibit generalization when unrepresentative samples are at hand—and it is the joint working of this package that constitutes human epistemic competence for induction. Common evaluative judgments regarding concrete episodes of belief-fixation reflect the requirement that the agent has developed and deploys the full package of processes that jointly constitute the relevant portion of epistemic competence. If the agent fails to either have or to deploy a process that is a part of the package making for epistemic competence, the agent "misses an important epistemic trick." The agent thus *comes to employ processes that are jointly less (globlally and transglobally) reliable than need be.*[7] Humans, normal adult humans, can do significantly better.

My argument in what follows will be that tractible background processes that serve to monitor interlocutors for trustworthiness are necessary to epistemic competence and must be in play in that processing by which a (normal adult) human agent arrives at an objectively justified belief by the reception of testimony. Such processes are required as complements to those that interpret one's interlocutor and that form the belief "on the interlocutor's say so." Such processes must be in play, at least in the background, just as (perhaps background) processes making for the sensitivity to sample characteristics must be in play in the competent and justified generalization from a sample.

My reasons for thinking that the objectively justifying processes of epistemically competent agents will include as a component processes that accommodate information bearing on interlocultor trustworthiness and regulate the acceptance of testimony accordingly, will include reliabilist themes such as those just expressed.[8] But, even before making this case, I wish to address an important issue having to do with the relation between episodes in which an agent is or is not justified in a belief and the idea of epistemic competence—and I wish to do so without relying on reliabilist themes. Consider a kind of case not involving testimony that instructively parallels testimonial cases. Suppose in a given case that an agent has a sample and has no information that would have called in question its representativeness. One

might wonder whether this agent needs to have in play background processes checking for representativeness—since, in just such a case, these processes would not make any difference to the results. Of course, there is this testimonial parallel: suppose in a case that an agent receives intelligible testimony to the effect that their car is due for an oil change, and that, as a matter of fact, the agent possesses no information that suggests the interlocutor here is untrustworthy. One might wonder whether this agent needs to have background processes in play that are sensitive to whether the agent possesses such information. I will soon argue that such background processes are required for competence in the case of testimonial belief formation, and for present purposes, I am flatly assuming (plausibly) that sensitivity to sample bais is required for epistemic competent generalization from samples. Right now, I want to address a general question: why is it required for objectively justified belief *in a given case* that the agent have, and have in play, the full complement of the processes that jointly make for epistemic competence. The question might be put this way: in a given episode or case, where a fully competent agent with just the information that the agent now possesses would arrive at the belief in question, can an agent be "episodically" objectively justified in that belief when not all the processes making for epistemic competence are possessed and episodically in play?[9] On the face of it, the answer should be negative. On the face of it, an agent with no information suggesting that a given sample is nonrepresentative, but who lacks processes sensitive to representativeness (or who for some reason does not have such sensitive processes in play, perhaps due to distraction), does *not arrive at the belief in question in the correct way*, would have generated the belief even were there compelling reasons to thing the sample was biased, and is thus not objectively justified. The suggestion is that episodic justification must be understood in terms of the general character of the cognitive processes in play in that episode—their general adequacy to deal with cases, and not in terms of some episode which luckily for the agent presents few pitfalls. The idea is that episodic justification must be understood in terms of systematic justification.

The suggestions just advanced can be supported, and in a way that does not suppose a commitment to reliabilism, or any other general account of what qualities of given constellation of processes would make that set of processes a variant of human epistemic competence. (What makes for

competence might be might be reliability of some stripe, or it might be sensitivity to coherence, or it might be sensitivity to evidence. We can be indifferent for present purposes.) What is crucial is the very general distinction between being objectively justified and merely having justification. We are concerned with being justified—and this requires that the belief in question be formed or maintained in the epistemically correct way. Famously, one can have justification for a belief and not be objectively justified or warranted in believing it. An agent, *A,* can form the exact same belief that a competent agent would have formed given the information *A* possesses,[10] and yet not have formed the belief in the right way. Perhaps *A* formed the belief by wishful thinking. That does not make for *A's* being objectively justified in the belief. Or, more to our point, perhaps *A* formed the belief *insensitively*—not in a way responsive to what information *A* possessed. Such insensitive belief fixation does not make for *A's* being objectively justified in the resulting belief. Note: this is so whether you think that the sensitive processes demanded are those that make for reliability, tracking of coherence, tracking evidential support, or some other matter that makes for good and justificatory constellations of processes. The significance of *sensitivity* here points beyond the particulars of the episode to counterfactual possibilities—to whether the agent's processes were of a sort that is "up to the tasks" presented by such cognitive contexts (or at least as up to the tasks as is humanly reasonable and tractible). This is why the lucky inductive generalizer is not justified—is neither episodically nor systematically justified—even when an inductively competent agent would there have arrived at the same result from the same sample and background information.

3. The Debate

I begin framing the issue by setting out two somewhat stylized positions.

3.1: Acceptance with Monitoring or Acceptance with Reason Inhibition

Some philosophers clearly hold that, in order for one to be warranted in believing on the basis of a piece of testimony, one must employ cognitive processes that monitor the trustworthiness of one's interlocutor or testimonial source. Call this the *Acceptance with Monitoring* (or AM) view. Fricker (1987,

1994) holds such a view, as do Adler (1994), Lyons (1997), and Falkner (2000), for example. I here support a version of the AM view.

Any plausible AM account will need to avoid unreasonable demands for monitoring processes. An adequate account of one's epistemic competence with respect to testimony will then need to indicate by what range of cognitive processes monitoring is to be tractably managed.

On the other hand, some philosophers find talk of "monitoring" troubling. Their basic picture is that one is entitled to accept the testimony one receives absent significant reasons for doubting the trustworthiness of its source. On this view, it is a fitting epistemic *default position* that one accepts testimony, that one simply trusts one's interlocutor. To say that trust is *a default* here signals (a) that one *can* have reasons for doubting the trustworthiness of one's interlocutor, (b) that, *when* one has such a reason, one then cannot simply accept the testimony (the default may give way, so that warranted acceptance depends on coming to terms with the apparent underminers), and (c) that absent such reasons, simple acceptance is warranted. (Such a view is championed by Burge 1993, 1997, and is congenial to many others, notably Coady 1992, 1994). Call this the *Acceptance with Reason Inhibition* view (ARI).

What is not generally appreciated, or well marked, is this: in keeping with the defeasible entitlement structure they describe, proponents of ARI must be read as supposing that there should always be complementary processes checking and regulating the acceptance of testimony—ones that are sensitive to reasons to doubt the interlocutor's trustworthiness. (This is the only charitable way of reading proponents of ARI.) In keeping with the *defeasible* entitlement structure that defines ARI (in particular, in keeping with clause (b) above), *were an agent to have some reason for doubting an interlocutor's trustworthiness, that agent would not be warranted or objectively justified in simply forming the indicated belief on the basis of that testimony—to be objectively justified in forming the testimonial belief, that agent would need to have in play processes that somehow came to terms with the possessed information bearing on trustworthiness.* Call this the *defeated entitlement case requirement.* When the agent has information suggesting interlocutor untrustwothiness, processes coming to terms with such information are required to be in play to yield an objectively justified belief. (So the proponent of ARI insists.) But,

imagine a position according to which one can, in other cases, form an objectively justified belief by simply accepting testimony, where this is unchecked or unregulated in any way by complementary processes. This would be to allow for unregulated acceptance of testimony in what might be termed a *bare satisfied entitlement case.* ARI cannot be charitably understood as coherently advancing such different demands (for objectively justificatory cognitive processes) in the two kinds of cases. Consider what is required of an agent's cognitive processes in the defeated entitlement case. In a defeated entitlement case, to be objectively justified in one's testimonial belief, the agent must have in play processes that are sensitive to the presence of information bearing on interlocutor trustworthiness. *Without such processes, the agent could not initiate or undergo processes that come to terms with such information when present—and this is what is required of the agent* in defeated entitlement cases. (This is not to say that the processes here required need be conscious, articulate, or at the level of internalist epistemology.)[11] To say that in such cases the agent *ought to* come to terms with the possessed information implies that the agent *can* come to terms with such information in such cases. Such *a capacity* on the part of the agent requires that the agent have or undergo *processes that are sensitive* to the presence of such information. So, *suppose that an agent has the capacities or functional processes required by ARI in connection with the reception of testimony in defeated entitlement cases. Of such an agent, with such capacities*, in a bare satisfied entitlement case, this must be true: *were that agent to have reason to doubt the trustworthiness of the interlocutor, that agent would initiate or undergo processes coming to terms with such information, and would not simply trust the interlocutor.[12] This much is true of any agent who has the capacity to satisfy the requirements of ARI in the defeated entitlement case—even when that agent happens to face a satisfied entitlement case.*

Thus, the proponent of ARI cannot simply distinguish two classes of cases (the bare satisfied entitlement cases and the defeated entitlement cases) and demand that (in an unconditioned manner) different processes be in play in the two. What is demanded of agents in the defeated entitlement case is that certain processes there be triggered. To demand this is to demand selectively triggering processes.[13] It is to demand that the agent would not simply undergo trusting acceptance in the defeated entitlement

case—and this is to require that the agent have in play (at least background) processes that are sensitive to the presence and strength of information bearing on interlocutor trustworthiness.[14] This is to require that the difference in the course of cognitive processing in the two cases be conditioned by the differences in the two cases—that processes in the agent be sensitive to which sort of case obtains, that even in the satisfied entitlement case the agent's background processes be sensitive to information bearing on trustworthiness of interlocutors.

One can read these last paragraphs as exhibiting how reflection on the defeasible entitlement structure envisioned in ARI leads one to think in terms of sensitive background processes, or constellations of yoked processes, and thus of the demands of the epistemic competence that is required for being objectively justified or warranted rather than merely having justification.

Proponents of ARI have been most articulate in characterizing what is distinctively required of the episodic processes in the satisfied default cases—and I am not here suggesting that they erred in so doing. Instead, I am prompted by what I think is rightly suggested by formulations such as Burge's *Acceptance Principle*: "*A person is entitled to accept as true something that is presented as true and that is intelligible to him, unless there are stronger reasons not to do so*" (1993, p. 467, italics in the original). I am suggesting that the "unless" clause here must be charitably understood as pointing to the second class of episodes—the defeated entitlement episodes—and to a background capacity to deploy different episodic processes in ways sensitive to these different cases. The distinctive episodic processes that ARI allows as justificatory in satisfied entitlement cases must be understood fitting when restricted as the unless-clause indicates. The episodic processes that ARI allows as justificatory in satisfied entitlement cases must be understood as episodically *fitted*—in terms of a systematic competence that includes background sensitivity to interlocutor trustworthiness.

ARI and AM may well be compatible, depending on how central elements of each are understood. With respect to AM, there is the crucial question of how the required monitoring is to be understood. With respect to ARI, there is the crucial question just how the complementary presupposed background processes are understood. With a demanding enough understanding of presupposed

background processes, and with a modest enough understanding of monitoring, ARI and AM could well be two faces of a single epistemically appropriate model.

3.2. Variations

We can distinguish more demanding and less demanding variants of both ARI and AM. Consider ARI. A relatively undemanding variant would make relatively weak demands on the background processes that are sensitive to reasons for doubting the trustworthiness of one's interlocutor. Call this view *Acceptance with Weak Reason Inhibition,* or AwRI—and let its distinguishing characteristic be that it does not require supporting cognitive processes that are themselves *educated and proactively sensitive* to indicators of the absence of trustworthiness. AwRI does not require that the agent employ processes that are sensitive to external clues regarding interlocutor trustworthiness. It does not require that the agent be sensitive to considerations bearing on the trustworthiness of the interlocutor, even when such considerations might seem obvious to most normal human observers. Instead, AwRI only requires that the agent be inhibited in accepting testimony by whatever *internal reasons* that agent *happens to have* for suspecting untrustworthiness, where the agent may be uneducated in the commonly appreciated marks of untrustworthiness, insensitive to such marks by virtue of never having garnered the background information commonly and readily acquired in their social context, and oblivious to such matters. At least with respect to AwRI, Fricker's charges would seem to have real application: an agent that is so uneducated and insensitive to real world indicators of untrustworthiness would possess few internalist reasons for inhibiting their acceptance of testimony and indeed would be gullible—"easily gulled"— despite the background processes envisioned. While the background processes are sensitive to what information the agent possesses regarding trustworthiness—it allows that they may be what would be— for normal human adults—"willfully" or negligently ill-informed on such matters. Such an agent seems epistemically unacceptable—and to yield objectively unjustified beliefs. Thus, an agent could satisfy AwRI and be unacceptably gullible.[15]

More plausible variants of ARI would demand more in the way of complementary background processes, and would consequently would be immune to Fricker's misgiving. Such an account of epistemic competence, call it *Acceptance with Strong Reason Inhibition* or AsRI, would require not just that there are in place processes appropriate to somehow dredging out of the agent's antecedent information, items (if any happen to be had) giving cause for doubting a given interlocutor's trustworthiness, but also that the agent have funded their memory with a reasonable sensitivity to such matters. This requires background processes that have come to be sensitive to the track record of various kinds of interlocutors. It requires that they include a reasonable acquired sensitivity to real world indicators of trustworthiness.[16] (It is worth noticing that nothing in AsRI requires the kind of articulate evaluation of an interlocutors trustworthiness that has worried and motivated proponents of ARI.).

Just as ARI can be understood as advancing an undemanding (AwRI) or a demanding (AsRI) understanding of objectively justificatory belief-fixing processes, so AM can be given stronger or weaker readings—depending on how one understands the demand for "monitoring." It is important that the demand be neither too weak nor too strong. The latter has been the more common worry, and most of my discussion will be devoted to avoiding that pitfall. But, it is worth noting at this juncture that, for an AM parallel to AwRI, one need only adopt a markedly undemanding understanding of what it is to monitor for trustworthiness.

Both AM and ARI accounts of the warranted formation of testimonial beliefs agree that epistemically fitting reception of testimony must be checked by processes that are sensitive to whatever reasons the agent happens to have for doubting the trustworthiness of the relevant interlocutor. That much is a minimum demand intrinsic to both; each requires sensitivity to undermining reasons, once those reasons are internally possessed by the agent receiving testimony. The AwRI variant of ARI is stipulated as making only this minimum demand—being satisfied with just sensitivity to reasons that happen to be possessed. In a parallel fashion, there is a possible variant on AM that requires of monitoring processes only that they be sensitive to whatever reasons for trusting or not trusting the interlocutor the agent happens to have. This variant of AM, call it *Acceptance with Weak Monitoring*, or AwM, would be

subject to the same misgivings that would plague AwRI. Stronger demands may be incorporated into AM or into ARI, and would involve greater demands on monitoring processes, or complementary background processes, had by an epistemically competent adult agent. This is to understand warranted testimonial belief fixation to require more than acceptance checked by possessed reasons; it is to insist that warrant comes only when agents employ processes that make for a a reasonable sensitivity to real world indicators of trustworthiness. Such a demand was incorporated in AsRI. A parallel variant of AM demands that monitoring for trustworthiness turn on processes with such acquired external sensitivities. Call this *Acceptance with Strong Monitoring*, or AsM. I believe that something like AsM has been the preferred position of many supporters of AM. AsM and AsRI are in accord in thinking of the epistemics quality of an episode of testimonially based belief formation as dependent on the epistemic quality of the range of epistemic processes there in play. If one's immediate processes are the result of being oblivious to what should be obvious to normal adult humans—where this results from a kind of epistemic negligence—then this colors the epistemic quality of one's testimonial reception in an episode. The idea that the epistemic quality of an episode, episodic justification, could be dependent on the epistemic quality of antecedent episodes and their traces is pervasive.[17]

Positions that call for monitoring as a precondition for warranted testimonial belief acceptance (at least in normal adults) might easily be conflated with the position commonly labeled *reductionist*. But, AM and reductionism are far from identical.

For the reductionist, either the acceptance of each piece of testimony must be individually given *nontestimonial* support or vindication—nontestimonial reasons for thinking it likely to be veridical—or there must be some sweeping argument for the conclusion that most testimony is veridical. In either case, the epistemic standing of testimony is reduced to the epistemic standing of the nontestimonial bases for such argumentation. One can readily envision a weaker sort of demand for monitoring that does not have this reductionist character—Nonreductive AM.[18] For reasons reflected in Coady's and Fricker's writings, I believe that real reductionism is a nonstarter. So I will largely ignore it in what follows.

Again, epistemic standards (our best understandings of objectively justificatory cognitive processes) are descriptively constrained idealizations characterizing a select set of tractable processes by which cognizers of the relevant sort can effectively pursue the production and maintenance of systems of true beliefs. [19] Such standards are revisable and evolving attempts to capture what actually constitutes the effective pursuit of understanding by systems of the relevant sort. Such an idealization provides an account of epistemic competence. Processes conduce to managing well our epistemic chores in yoke with other processes. Thus, epistemically justificatory processes comprise a package whose *joint* realization or implementation conduces to epistemic success. The range of processes that make for human epistemic competence, and whose deployment is requisite for the generation or maintenance of objectively justified belief, comprise a set of mutually complementary (sometimes checking or restraining, and sometimes augmenting or extending) processes jointly contributing to the successful management of the agent's epistemic chores.

Consider that core process that implements what one might term the *simple acceptance* of testimonial beliefs. This process takes as input the output of those processes that make for competent interpretation of interlocutors. [20] The interpretive processes provide the recognition of an interlocutor (or source) as having asserted something, *P*. To a crude first approximation, one can then think of the simple acceptance process as responsible for the consequent formation of the belief that *P* "on that person's or source's say so." Whatever the more detailed account one gives of this process, all parties (all variants on ARI and on AM) would see something along these lines as a core component of normal adult epistemic competence with testimony. Their differences are over what processes must be yoked with the core processes of simple acceptance to make for normal adult human epistemic competence and objectively justified belief formation.

As we have already seen, all variants of ARI, even the least demanding, require that simple acceptance be supplemented by processes that check or restrain it. Without such restraining processes the formation of testimonial beliefs would not have the *defeasible*-default character envisioned and required by AwRI (and by ARI generally). AwRI requires processes that restrain acceptance when the agent

happens to have reasons for doubting the trustworthiness of the interlocutor, but does not require that these processes themselves be such as would lay in a reasonable stock of sensitivities to external indicators regarding trustworthiness. AsRI requires more just where AwRI is tolerant.

As we have also seen, approaches that demand that the agent (nonreductively) monitor for the trustworthiness of the interlocutor or source, indeed all variants of AM, demand processes that modulate or regulate the process of simple acceptance. The commonly preferred variant would be AsM, which requires regulating processes that *are sensitive to the trustworthiness of interlocutors or sources*. But, at least in logical space, there are variants of AM that parallel the possible variants of ARI, demanding more or less of the monitoring processes that modulate simple acceptance.

Readers should be wondering whether there is any difference between ARI and AM. After all, it turns out that each supposes that some process of simple acceptance needs to be modulated or regulated by background processes. For any variant of the one, there would seem to be a variant of the other that is characterized by making parallel demands on the complementary regulating processes. There does seem to be at least this difference: ARI envisions processes that regulate simple acceptance by inhibitinig it, whereas AM envisions processes that may regulate simple acceptance in additional ways—say by inducing it as well as by inhibiting it. Further, on AM accounts, it may be that simple acceptance does not come into play without being induced by background processes (at least in mature human cognizers).

One question that can be posed regarding both the complementary background processes of ARI, and of the regulating monitoring processes of AM, concerns whether these are supposed to proceed at the level of the classically accessible, as the traditional access internalist might suppose. Indeed, the issue of internalism and externalism are never too far from the surface in the testimony literature. For example, Burge's *Acceptance Principle* enunciates a classic formulation of what I am calling ARI: "*A person is entitled to accept as true something that is presented as true and that is intelligible to him, unless there are stronger reasons not to do so*" (1993, p. 467, italics in the original). Here, talk of reasons might be understood as suggesting an internalist account, but this would be mistaken. Burge's talk of "entitlement" is used to indicate that full internalistic vindication is not necessary for the warrant—'entitlement' signals

an externalist dimension to warrant. For our purposes here, this suggests that background complementary processes might fruitfully be taken up from the perspective of an externalist epistemology. At the same time, one senses that misgivings regarding monitoring requirements sometimes have to do with whether it is plausible that some kind of access internalist process must license all instances of acceptance. That seems clearly intractable, and such resistance seems warranted.

There really are two issues here, and I want to focus mostly on the second. First there is the question of whether the acquired information about, and sensitivity to, variations in trustworthiness (that for which background processes might be held responsible in AsRI and in AsM) is generated by processes that are of a sort that an internalist can love. This will not be my central concern. Second, there is the question of whether the monitoring or regulating processes that draw on such information, those processes that rather directly regulate acceptance, are to be thought of as turning on accessible resources, accessible considerations. Are these to be held to conform to internalist strictures so that they can be understood as a matter of articulately weighing reasons for thinking trustworthy or untrustworthy? The idea that the monitoring required for objectively justified testimonial belief formation would need to be a matter of articulate, or at least fully articulable, determination of the likely trustworthiness of one's interlocutors, and that this is required by AM, may account for some of the misgivings about AM. I will soon argue that monitoring should not be so understood. But, it is worth noting that some ARI thought may be motivated by the recognition that such an internalist monitoring requirement would be unworkable.

It is noteworthy that Coady, a writer prominently associated with the Neo-Reidian, ARI, accounts of testimony, seems to envision the normatively appropriate reception of testimony as involving an important complementary background sensitivity on the part of agents that seems largely managed at a level that would elude full internalist analysis. Thus, he remarks:

What happens characteristically in the reception of testimony is that the audience operates a sort of learning mechanism which has certain critical capacities built into it. The mechanism may be thought of as partially innate, though modified by experience, especially in the matter of critical

capacities. It is useful to invoke the model of a mechanism here since *the reception of testimony is normally unreflective but is not thereby uncritical*. We may have 'no reason to doubt' another's communication even where there is no question of our being gullible; we may simply recognize that the standard warning signs of deceit, confusion, or mistake are not present. This recognition incorporates our knowledge of the witnesses competence, of the circumstances surrounding his utterance, of his honesty, of the consistency of the parts of his testimony, and its relation to what others have said, or not said, on the matter (Coady 1992, p. 47, italics added).

It would be hard to improve on these remarks as a call for background processes that only partially ascend to the level of internalist accessibility.

Fricker, whose AM account reflects a general internalist/coherentist approach to epistemological matters, nevertheless also seems to envision monitoring processes that extend beyond what is classically accessible. On her account, "the hearer must always be monitoring the speaker critically, where this monitoring is a matter of *the actual engagement of a counterfactual sensitivity*: it is true throughout of the hearer that if there were any signs of untrustworthiness, she would pick them up" (1994, 154; italics added). Of course, such "counterfactual sensitivities" are not the usual elements of internalist accounts. Fricker's formulation (that it should be "true throughout of the hearer that if there were signs of untrustworthiness, she would register them, and respond appropriately," Fricker 1994: 150) suggests an understanding of monitoring according to which it gives rise to internalist moments, perhaps as a part of the immediate gatekeeping for acceptance, while much of the real substantive, sensitive, work is managed in ways that would escape internalist description. As she seems to recognize, to require that more be done at the level of internalist radar would make a particularly demanding, overly demanding, account of warranted testimonial belief formation.

So, as a final element of this survey of possible ARI and AM positions, we might want to include mention of versions of AsRI and AsM that would envision the relevant background or monitoring processes being managed wholly at the level of the classically internalistically accessible. Call these

AsiRI and AsiM. One benefit of distinguishing such positions is that it makes clear that such a thoroughly

internalist accessibility demand need not be made by the more generic AsRI and AsM.

Bringing all this together, we might provide a rough taxology of the approaches to the

epistemology of testimony (figure 1):

Figure 1

ARI—The epistemic *default* is that one can simply accept testimony. Here, (a) one *can* have reasons for doubting the trustworthiness of one's interlocutor, (b) *when* one has such a reason, one then *cannot* simply accept the testimony, and (c) that absent such reasons, simple credulous acceptance is warranted. Variations involve different demands on background regulating processes.	

AwRI:
    a. requires processes appropriate to dredging out of the agent's antecedent information, items (if any happen to be had) giving cause for doubting a given interlocutor's trustworthiness,
    b. does **not** require supporting cognitive processes that are educated and proactively sensitive to indicators of the absence of trustworthiness. Allows these processes to be "stupid" and obtuse.
	AsRI:
	    a. requires processes appropriate to dredging out of the agent's antecedent information, items (if any happen to be had) giving cause for doubting a given interlocutor's trustworthiness,
	    b. *additionally* requires supporting cognitive processes that are educated and proactively sensitive to indicators of the absence of trustworthiness.

AwiRI—this internalist variant of AwRI demands that the background regulating processes operate on internally accessible, possessed reasons.	**AsiRI**—this internalist variant of AsRI requires that the background regulating processes be managed at the fully articulate or articulatable level of argumentative vindication of interlocutors.

Nonreductive-AM—Testimonial belief formation must be regulated (encouraged, allowed, and inhibited as fitting by processes that monitor interlocutors for trustworthiness.	

AwM—
    a. requires monitoring processes appropriate to dredging out of the agent's antecedent information, items (if any happen to be had) relevant to interlocutor's trustworthiness,
    b. does *not* require monitoring processes that are educated and proactively sensitive to indicators of the absence of trustworthiness. Allows these processes to be "stupid" and obtuse.
	AsM—
	    a. requires monitoring processes appropriate to dredging out of the agent's antecedent information, items (if any happen to be had) relevant to interlocutor's trustworthiness,
	    b. *additionally* requires background processes that are and have been sensitive to the track record of various kinds of interlocutors.

AwiM—this internalist variant of AwM demands the objectively justificatory monitoring processes be wholly at the level of what is classically accessible.	**AsiM**—this internalist variant of AsM demands the objectively justificatory monitoring processes be wholly at the level of what is classically accessible.

Reductive-AM—Testimonial belief formation must be regulated (allowed and inhibited, as fitting) by processes that monitor interlocutors for trustworthiness—where those monitoring processes ultimately rest only on sources that are not testimonial in character.
    This position would be subject to variations parallel to those represented above.
    These are treated as nonstarters, for reasons familiar in the literature	

4. Strong Regulated Acceptance, Why it is Not Optional

4.1: A Commonsense View of the Monitoring or Complementary Processes Called For

Reflection on commonsensical epistemic demands seems to support a position in the neighborhood of AsRI or AsM. In the course of growing up and living in a human society, all of us have had rather a lot of experience with various sources attesting to various matters. A lot that one has come to understand about the world has come at least partially in this manner. Sometimes we have been pleased with the resulting understandings, in other cases we have been disappointed. In the bargain, a normal adult human will have *acquired significant stores of information and important sensitivities about testimony* itself. In some crude and commonly inarticulate fashion, one should and normally will have come to understand rather a lot about testimony. One who has not is, in an important epistemic sense, either defective (because of cognitive defects that are atypical of humans) or remiss. So, for normal adult humans, one who has not acquired some significant fund of information and sensitivities would be judged remiss—as failing to develop and employ an acceptable epistemic competence—and would then be judged unwarranted or not objectively justified in beliefs that diverge from what would have been formed were such a competence in place.

One will have some understanding of how various interests may affect the tendency to tell the truth, the whole truth that would seem relevant to a matter under discussion, or a select set of truths, or straight falsehoods. One will have come to understand various social roles, and to understand the range of interests (reward structures, legal liabilities, and so forth) that may be commonly possessed by folk occupying such roles. (Some of these understandings have come at least partly by way of testimony.) One has developed rather sensitive recognitional capacities that we employ and refine in our everyday interactions. For example, some of these capacities make for sensitivity to facial gestures and "body language." Presumably one also has some appreciation for how these things can be manipulated. Reader's who wish to exempt themselves from these characterizations may do so. It is admittedly easy to

overestimate one's capacities. But, most of us have attempted to cultivate these capacities as far as possible—and *most find cultivation to a reasonable degree to be epistemically fitting and a non-optional component of epistemic practice*.

To the extent that one is uncertain of one's own capacities to detect deception on the part of those who might be interested in deceiving, one often attempts to limit one's reliance on particular pieces of testimony—one looks for information from independent sources, some of which may be "more trustworthy." (For example, one may look for pricing information from consumer-oriented publications before shopping for a used auto—even when the sales staff assure us they are moving their commodities at a loss. One may have an independent technician look over a candidate auto for hidden defects—despite assurances that it is in mint condition.) One relies on alternative sources—for example, interlocutors that make their living on the basis of iterated episodes of providing information on which various others may eventually form opinions regarding truthfulness. Professional reputations are at stake, and one is not blind to the social mechanisms that help to motivate trustworthiness. The commonsensical bottom line: *One expects that normal adult humans will have had a representative range of experience with interlocutors and have acquired information, capacities, and practices such as those just mentioned*.[21]

An interlocutor's tendency to tell what is believed to be the truth is only part of what makes for trustworthiness. Some matters are difficult to ascertain even with the best of intention, and one comes to discount testimony on these matters accordingly. Few of us are lacking in earnest want-to-be spiritual advisors, for example. Most of us form expectations regarding such advisors. (One presumably has comparatively higher expectations for the information in elementary chemistry or biology books.)

One could elaborate these crude observations, and such elaboration would elicit and reflect commonsensical expectations and standards for what it is to be a *competent epistemic agent*, and what is required for *epistemically warranting testimonial belief fixation*. This elaboration would highlight the kinds and amount of information that competent agents would have come to possess regarding testimony. Indeed, and this is important, *one would judge epistemically defective or extraordinarily remiss an adult in any culture who had not acquired much such information and much such sensitivity to testimony*—one

would rightly think of such an agent as a kind of epistemic disaster waiting to happen. To the extent that an agent lacked such information and sensitivity, one would not trust them as a full-fledged member of one's epistemic community. Such an agent would lack a component of epistemic competence that we all find highly valuable or desirable—and one evaluates the agent's epistemic practice and character accordingly. In so doing, one mitigates the threat that such simple credulous souls will pollute one's epistemic community by accepting and passing along problematic beliefs accepted too freely from sources that "ought to have been seen to be suspect." An agent who could not acquire such information and sensitivity in the normal course of growing up would perhaps be considered a kind of unfortunate perpetual child, to whom our standards of normal adult epistemic competence do not apply. But a normal agent who could, but does not, acquire and employ such information would be judged more harshly.

What all this points to is that our epistemic standards are such that we demand of agents not just that they forebear from accepting testimony when they happen to have reasons for thinking it untrustworthy, but also that agents be *informed and sensitive monitors of conditions and characters that foster or compromise trustworthiness*. One seeks to cultivate such information and sensitivity in oneself, and one looks for it in others. It is an integral part of our epistemic evaluative practice that we demand of agents (ourselves and others) that they be *suitably trained up monitors of trustworthiness*. This is treated as a *non-optional component of competent, objectively appropriate, epistemic practice* in connection with testimony. In this section, I have sought merely to highlight a very commonsensical expectation for normal adult humans within any community, and to thereby make somewhat articulate a not at all shocking demand for a kind of sensitive monitoring of interlocutors for trustworthiness.

What I have just said is not intended to distinguish between AsRI and AsM; both require that the competent agent have developed and have in play processes that are informed about, and sensitive to, indicators of trustworthiness. Both require that the process of simple acceptance be regulated by such processes. This is a good start. But, it is important to show that the suggested commonsensical demand is not merely commonsense, rather, there is a general epistemological basis for such demands. Ultimately,

doing so revolves around the question of what package of humanly tractable cognitive processes has a satisfactory balance of global or transglobal reliability, power, and speed.

4.2: Strong Regulated Acceptance

One component process will need to be included in any understanding of the epistemology of testimonial beliefs; it is that core process that realizes what was earlier termed *simple acceptance*.[22] This process takes as input the interpretation of interlocutors as making an assertion. It is then responsible for the consequent formation of a belief with the asserted content "on that person's or source's say so." So understood, it is a relatively simple process. Burge (1993, 1997) compares the acceptance of testimony to the formation of beliefs on the basis of memory. In the case of memory, one recalls an episode and straightway forms the belief that things are as there recalled. For example, one recalls having derived a line in a proof, and then, one simply proceeds with what was derived there. Or, one recalls what one did as one entered the house, and forms the belief that one placed one's keys on the kitchen table. The suggestion is then that acceptance is a similarly simple or direct process.

The acceptance of testimony has also been compared to the formation of beliefs on the basis of perception. Again, one point of such comparisons is clear: insofar as perceptual beliefs arise simply or directly out of perceptual experience, acceptance of testimony is conceived as a process in which testimonial beliefs arise directly out of interpretive recognition of what one's interlocutor asserts. Here the idea is that the proper epistemic processes of forming perceptual beliefs, given one's perceptual experience, do not need special permissions, and that one can and should simply accept the belief reflecting the perceptual seeming. One idea behind insisting on a parallel between the simple acceptance of testimonial beliefs and the formation of perceptual beliefs is that acceptance can, epistemically appropriately, be a straightforward matter of forming a belief with the same content as that testimony that one's interlocutor is interpreted as having provided—that the transition from interpreted testimony to belief formation can be as simple and unproblematic as the transition from a perceptual seeming to the

indicated perceptual belief. Acceptance takes as input the interpretation of the interlocutor's utterances and simply forms a belief with the corresponding content (unless inhibited).

These comparisons with memory and perception will be distorting, unless they are understood in terms of two further points.

First, as discussed already, it is epistemically undesirable that simple acceptance be unchecked, unmodulated, or unregulated in any fashion in any instance. As I have interpreted the debate, this much is really not in contention. Both ARI and AM (even in their weakest variants), require that there always be processes in place within each agent that are sensitive to whatever accessible information that agent happens to possess in order to turn up combinations of information that would count as a reason to doubt trustworthiness. Both insist that such information must be accommodated in the course of testimonial belief formation. With fitting background processes in play, the transition from interpreted testimony to belief may be smooth and direct at the level of articulate processes. At the same time, humans can have (rather ill-understood) background processes by which they automatically "survey" their standing information for that bearing on the question of concern.

Second, the suggested comparison with perceptual competence is instructive only when one thinks realistically about what makes for perceptual competence. What is it to be perceptually competent with respect to some limited matter? What is it to be perceptually competent generally? When one is a competent perceptual judge of the wildlife in one's environment, for example, one can with reasonable sensitivity respond to common episodic encounters (and to more or less enduring traces) in one's environment. Due to one's sensitivity, one can produce "reasonably" accurate identifications of the species of wildlife involved in the episodes (or leaving the traces). One who is perceptually competent with respect to a subject is able to make reliable judgments in a range of common environmental conditions. Thus, one who is perceptually competent with respect to local wildlife can make reliable identifications when various creatures are partially obscured, quickly glanced, or imperfectly illuminated (common conditions). Importantly, such capacities in turn require reasonable sensitivity to when conditions do not allow reliable identification. *A competent perceptual judge on a given subject matter is*

thus one who has been trained up to be reasonably reliable on the matter in question—one who can consequently render verdicts that are likely true—and this requires not rendering verdicts in certain ranges of epistemically difficult cases.[23]

What then of an agent who is perceptually competent generally? This would seem to be an agent who has acquired perceptual competence on some reasonable range of subject matters, and who tends to refrain from forming perceptual judgments on subject matters where such competence is lacking.

When thinking about the acceptance of testimony on the model of the formation of perceptual beliefs, it is important to not loose sight of the richness of our perceptual processes. It has sometimes seemed tempting to think of perceptual processes as epistemically simple black boxes. Such understandings of the processes involved in perception are themselves epistemologically inadequate. Let us explore the parallel that seems suggested by comparing perceptual belief formation and testimonial belief formation. It is important to recognize that a rich set of acquired information and sensitivity informs both the generation of perceptions themselves and the subsequent smooth or hesitating subsequent move to perception beliefs. In parallel fashion, we will want to think of both the generation of interpretive understandings (the crude parallel to the generation of perceptions) and the subsequent smooth or hesitant generation of testimonial beliefs as themselves conditioned by rich sets of acquired information. The interpretive understanding of an interlocutor as asserting some proposition stands as input to testimonial belief formation, and this is reasonably thought to be parallel to how a kind of perceptual seeming is to serve as a basis for the formation of a perceptual belief. To make this parallel plausible, it is important that the perceptual seeming be thought of as reasonably rich and substantive. Suppose that I form a perceptual belief that there is a snake there in the sun at the side of the trail. If this is to arise fairly straightforwardly out of the perceptual input, then the perceptual input must not be caricatured, along the lines of some epistemologist's cartoon, as a pattern of simple qualia. The perceptual input must be a kind of substantive perceptual seeming—the appearance a snake there in the sun, or something in this ballpark.[24] Let us call these rich perceptual inputs to perceptual belief formation *perceptions*. Whatever goes on in the processes yielding such perceptions must itself be quite

sophisticated and complex. If the epistemically competent formation of perceptual beliefs is to serve as a helpful model for thinking about epistemic competence in the formation of testimonial beliefs, we must not neglect this complexity. It will be necessary to be alive to the ways in which epistemic competence requires rich informed sensitivities that play a role in the formation of perceptual beliefs. I suggest that epistemically desirable perceptual processes turn on sophisticated trained sensitivity, that such sensitivity involves the automatic accommodation of much information that the system somehow has come to possess from previous training, and that consequently competent perceptual processes always already involve a crucial kind of regulation. One's perceptual processes (both the perceptions and the move to beliefs) are sensitive to light levels, shadows, degrees and placement of occlusions, context, and the like. They are shaped by past feedback and have come to accommodate rich stores of information automatically, so as to exhibit a "counterfactual sensitivity" to such matters. These ideas find expression in Henderson (1994) and in Henderson and Horgan (2000a).

What makes for competence in the formation of testimonial beliefs should be seen as parallel to what makes for perceptual competence: one who is competent in forming testimonial beliefs can respond to testimony from the sort of interlocutors commonly encountered in one's environment by forming beliefs that have a reasonable chance of being true. Given the diversity in interlocutors, this requires significant discriminative sensitivity—a sensitivity that must be acquired by training and by coming to understand social roles (as perceptual sensitivity must be acquired). Such an agent will form testimonial beliefs in certain cases and forebear in others, and in each sort of case there is in play an informed capacity making for sensitivity to environmental cues reflecting trustworthiness. One of normal human capacities who yet has somehow neglected such matters is epistemically both deficient and remiss—and, as a result of the resulting incompetence, such an agent is not warranted in forming beliefs by testimony. An adult who has remained in so ill-informed and insensitive to testimony, and who yet forms a belief on the basis of testimony, is like one who, having neglected perceptual training in some matter (say a North American who has no familiarity with the continent's snakes), yet confidently forms a perceptual belief

that yonder is one of those things (say a cottonmouth). Such are not objectively justified or warranted beliefs.

If the acceptance of testimony on the basis of interpretive recognition is understood as like the formulation of perceptual beliefs on the basis of perceptions, it would then seem correct to think of the interpretative understandings as standing in the course of testimonial belief formation as the parallel to perceptions that (smoothly or hesitantly) foster the formation of perceptual beliefs. I suggested that perceptions themselves are sensitive to, or conditioned by, rich sets of antecedently possessed information—such is perceptual competence. It is certainly significant that interpretive understandings are themselves sensitive to much antecedently possessed information. Interpretation itself always turns partly on issues of motivation, likely error, and the like (Davidson 1984, Henderson 1993). Consider a simple example. Perhaps during the presidency of George W. Bush, one happens to hear a person on television discussing the "elegant and perspicuous case developed by the President." Before one arrives at an interpretive understanding of the commentator and their utterance, one probably looks for cues. Do you recognize the commentator as a Republican mouthpiece? As a comedian? Are they smirking? Do they look drunk? Are they at a NASCAR event, or on a hunting and fishing show? Are they associated with an energy company? Is there a church with a field of small white crosses and anti-abortion placards in their background? All this (and much much more) can color the upshot to the interpretive question: What really is being said, and why? Clearly, interpretation is subtle process that always already is sensitive to rich sets of information relevant to trustworthiness. The regulation of the formation of testimonial belief is already underway in the process of forming an interpretive understanding.

To demand (with AsM or AsRI) that the process of epistemically competently forming beliefs on the basis of testimony be strongly regulated is to require that the competent agent have acquired a sensitivity to the trustworthiness of interlocutors and sources. Such an agent will be reasonably likely to have reasons for thinking that the interlocutor is not (or is) trustworthy when that interlocutor is not (or is) trustworthy—or, if not full-blown reasons, at least difficult to articulate bases or senses for such matters in some cases.[25]

4.3: In praise of strong regulated acceptance

Acceptance with strong monitoring or with strong reason inhibition has significant advantages over simple acceptance, and over acceptance with weak monitoring or weak reason inhibition. The advantages are matters of gains in global and transglobal reliability. As the result, monitoring is a non-optional component of epistemic competence for normal adult humans (provided it can be managed without incurring high costs in terms of power and speed). This is to say, just as one would judge inadequate, and not yielding of justificatory beliefs, inductive processes that are insensitive to sample bias, so one would judge that testimonial belief formation that is insensitive to interlocutor trustworthiness does not yield objectively justified belief.

i. Gains in Global Reliability

The real world is one in which interlocutors and sources differ markedly in trustworthiness. Interlocutors can fail to be trustworthy on a given matter as a consequence of one or a combination of two sources. Interlocutors can fail to believe the truth on a matter, or they can be motivated to not reveal what they know or to provide false information on a particular matter. So we live in a world with such variations in trustworthiness. Normal adult humans have also learned rather a lot about such variations in trustworthiness. To whatever extent one can effectively bring such information to the task of gauging particular pieces of testimony, one comes to attend selectively to more reliable testimony, and the global reliability of testimonial belief formation is generally improved. Again, as one discounts the testimony of (some of) those whose trustworthiness on the relevant matter is compromised, one comes to rely on testimony that is correspondingly more trustworthy as a whole. As a result, the global reliability of one's testimonial belief formation is improved, sometimes markedly improved. Monitoring global enhances reliability.[26]

ii. Gains in Transglobal Reliability

Now imagine a range of epistemically possible environments, and think of the training that humans would there undergo in the course of their everyday experience, and of the trained and strongly

related acceptance that would result. A simple point emerges, across a range of environments, drawing on a course of experience, the resulting sensitive regulated acceptance would enhance the reliability of testimonial belief forming processes to satisfactory levels in those environments. Consider some cases.

One can start by imagining an environment in which all agents are trustworthy. What a wonderful, benign, environment! Simple acceptance, without monitoring or background regulation, would be reasonably reliable there. Of course, some form of regulated acceptance would not hurt here, and it would be reasonably costless. After all, in the benign environment, trained sensitivity would yield a correspondingly uniform trusting acceptance.

Towards the other extreme, we can imagine various environments in which interlocutors and sources are frighteningly coordinated and systematic in deceiving some agent. In the 1998 movie, *The Truman Show*, the main character has been raised from childhood on a movie set where he has lived his entire life being the subject of a massive deception. One might add to this scenario more malicious actors coordinated in an effort to see just how systematically mistaken a belief system they could produce in the agent they deceive. This is the social analogue of a demon-infested environment, an environment in which the agent is from birth at the mercy of a committee of skilled deceivers, including deceiving interlocutors. Of course, here the normal human agent will have acquired a kind of sensitivity by training—a sensitivity to *apparent* trustworthiness. But, this will not yield significant rewards in terms of global reliability. As with demon-infested environments generally, so here there is little or no prospect for globally reliable belief fixation (if the committee is coordinated and skillful enough).[27]

However, between such extremes, there seem a range of environments that are rather like our actual environment, varying as only marginally better or worse in terms of the trustworthiness of testimony an agent is likely to encounter. In such environments (as well as trivially in the extremely benign ones), a normal human agent can (and naturally would) acquire a sensitivity to some significant portion of the variations in trustworthiness found there. What I have been calling the strong regulation of acceptance will then bring attendant benefits in the global reliability of testimonial belief formation in each of these environments. Thus, strongly regulated acceptance would be globally reliable in a wide

range of environments in which some process of testimonial belief formation would be reliable. Acceptance with strong regulation is transglobally reliable.

Thus, strong regulation of contributes to both global reliability and the transglobal reliability of testimonial belief formation. To the extent that normal humans can acquire processes for such sensitive or strongly regulated acceptance, to the extent that such processes are tractable and would be reasonably noncostly in terms of the overall power and speed of testimonial belief formation, such processes are rightly counted as a part of the epistemic competence of normal human adults for such belief fixation. When such sensitivity is not in play, the agent is "missing a trick," an important epistemic chore is omitted, and the resulting general beliefs are not objectively justified.

4.4. Tractable, Holistic, Sensitive Regulation

Reflection on commonsensical expectations for a reasonable sensitivity to trustworthiness on the part of normal adult human agents at home in their culture provides a prima facie case for thinking that the monitoring called for here is tractable, even common, in practice. I will here content myself with making a suggestion about how this can be.

In *Iceberg Epistemology*, Henderson and Horgan (2000a) argue that some epistemically important processes can differ markedly from the sort of processes that have commonly and traditionally been thought to exhaust those of epistemological concern. In particular, we argue that an epistemically competent finite cognitive system (such as a human), faces certain cognitive chores that require reliance on processes of sorts that do not readily fit into traditional epistemological accounts. One must recognize the epistemological importance of cognitive processes turning on information that is possessed by agents in ways that epistemologists have tended to overlook. Examples of the chores in question include the kind of holistic belief fixation reflected in theory choice, and the kind real time perceptual discrimination in the face of challenging conditions. These chores require that the cognitive system be able to somehow bring together the relevant information, sensitively gauging what information is relevant and taking account of that information. The problem for contemporary epistemologists is one of appreciating how this can be

done in ways that do not make for intractability. *Iceberg Epistemology*, is inspired by recent work in cognitive science and reflection on connectionist systems. It argues that a system can come to possess much information that is not represented in the system, and that the system can automatically accommodate such information without needing to represent it. I believe that these ideas at least reflect the direction in which epistemologists should look when attempting to understand not just theory choice and perception, but also the sort of ready way in which human being can be sensitive to many cues and much learned information when monitoring interlocutors for trustworthiness.

Fodor (1983) argues that holistic dimensions of central belief fixation make for insuperable problems for traditional computational cognitive science. These problems—the frame-type problems—arise largely because of the apparent computational intractability of managing all relevant information, insofar as that information gets *explicitly represented* in the course of cognitive processing.[28] What this suggests is that belief fixation and related cognitive processes operate in a way that accommodates much relevant information *automatically* and *implicitly*. The suggestion is that the holistic aspects of belief fixation involve not the finding and fetching of relevant representations from memory-banks where they are stored in explicit form, and not the overt representation and comparative evaluation of large-scale alternative belief-systems. Rather, these holistic aspects are somehow implicit in the structure of the cognitive system, in such a way that temporal transitions from one occurrent cognitive state to another accommodate the holistic aspects automatically. In the terminology of Horgan and Tienson (1995, 1996), the holistic informational content is *morphological*, rather than occurrent.

Morphological content is information that:

i)   is implicit in the standing structure of the cognitive system (rather than explicitly represented in the system's occurrent cognitive states or explicitly stored in memory), and

ii)   gets accommodated in cognitive processing without getting explicitly represented in occurrent cognitive states, either conscious or unconscious.

The apparent moral: in general, human belief fixation *must* operate in a way that draws heavily upon morphological content, in order to avoid computational intractability. One might say, these processes are *essentially* morphological.

This suggestion (developed further in Henderson and Horgan 2000a and forthcoming 2) offers a general, and not implausible, answer to the extremely daunting question of how the holistic aspects of belief fixation could be handled. Not unrelatedly, it provides a promising model for how both subtly informed perception and strongly regulated acceptance of testimony can be managed in tractable ways: In such matters, the competent epistemic agent will have acquired massive amounts of information and will accommodate that information in a largely automatic fashion.

5. Conclusions

Epistemically competent and justificatory belief formation on the basis of testimony will conform to something along the lines of AsRI or AsM. Both commonsensical expectations regarding the reception of testimony (section 4.1) and reflections on what package of cognitive processes would adequately further our epistemic ends (sections 4.2 and 4.3) point to the conclusion that warranted testimonial belief formation on the part of normal adult humans must be regulated by a rich acquired sensitivity to indicators of degrees of trustworthiness. Further, reflection on the sorts of cognitive processes that must serve in the context of holistic belief fixation and in perception, gave us reason to think that this sensitive regulation of testimonial belief formation also must be managed by processes that accommodate rich information automatically. In large measure, it must be managed essentially in the background (essentially morphologically) (section 4.4). This requires that one repudiate AsiRI and AsiM as unworkably demanding, concluding that these positions require that regulation be managed by processes that are intractable for creatures like us. Of course this is not to deny that there are significant internalistically accessible moments where background process flag important considerations, providing warnings or assurances that are then salient at the level of what is classically accessible.

There remains an issue that may have troubled readers for most of these pages. I have written of the "sensitive regulation" of the acceptance of testimony. This was to express a studied neutrality as to

whether AsRI or AsM provided the best account of the epistemology of testimony. The reason for reticence on this issue should be clear from section 3: AsRI and AsM must be understood as highly similar. There is simply "not much daylight between them." The range of variants of the one can be paralleled in variants of the other. The preferred variants, AsRI and AsM, are not as divergent as might have been supposed. Still, there would seem to be this difference: ARI envisions processes that regulate simple acceptance largely by inhibiting it, whereas AM envisions processes that may regulate simple acceptance in additional ways—say by inducing or encouraging it as well as by inhibiting it. (At least it would seem more open to this.) Further, on AM accounts, it may be that simple acceptance does not come into play without being induced by background processes (at least in mature human cognizers). I do think that these differences (particularly the first) provides a basis for preferring AsM over AsRI. Keeping in mind the character of the rich regulative sensitivity envisioned in this paper, it would seem strange were that sensitivity to play a largely inhibiting role as suggested in AsRI. For example, commonsensically, that sensitivity often plays a role in selecting those we seek as informants. Here it would seem that the sensitivity envisioned plays a role directing the obtaining of testimony, and provides at least a positive prima facie basis for its acceptance. Or consider again interpretation, which furnishes the input for testimonial belief formation. A realistic understanding of interpretation has it suffused with rich information, and from the inception engaged with indicators of trustworthiness. Only against such a background could one "immediately" perceive some utterance as subtle or not so subtle irony, an expression of helpfulness, perhaps an overture of friendship, a tolerant moment, an official chore, or the like. Thus, commonly, interpretation does not merely give a recognition of some simple content, but rather an appreciation of, or first take on, some actions thereby performed by one's interlocutor. It thus seems reasonable that such information provides sometimes encouragement and sometimes inhibition from the very start of testimonial reception.

References

Adler, J. (1994) "Testimony, Trust, and Knowing," *Journal of Philosophy* 1994: 264-275.

Audi, R. (1997) "The Place of Testimony in the Fabric of Knowledge and Justification," *American Philosophical Quarterly* 34: 404-417.

Burge, T. (1993) "Content Preservation," *Philosophical Review* 102: 457-88.

Burge, T. (1997) "Interlocution, Perception, and Memory," *Philosophical Studies* 86: 21-47.

Chernaik, C. (1986) *Minimal Rationality*, M.I.T. Press.

Coady, C. (1992) *Testimony: A Philosophical Study,* Clarendon Press, Oxford.

Coady, C. (1994) "Testimony, Observation and Autonomous Knowledge," in B. K. Matilal and A. Chakrabarti, eds., *Knowing from Words*. Kluwer Academic Publishers, pp. 225-250.

Davidson, D. (1994) *Inquiries into Truth and Interpretation.* Oxford: Clarendon Press,

Faulkner, P. (2000) "The Social Character of Testimonial Knowledge," *Journal of Philosophy* 97: 581-601.

Fodor, J. (1983) *The Modularity of Mind: An Essay in Faculty Psychology*, M.I.T. Press.

Fricker, E. (1987) "The Epistemology of Testimony," *Proceedings of the Aristotelian Society*, Supplemental Vol. 61: pp. 57-83.

Fricker, E. (1994) "Against Gullibility," in B. K. Matilal and A. Chakrabarti, eds., *Knowing from Words*. Kluwer Academic Publishers, pp. 125-61.

Fricker, E. (1995) "Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony," *Mind* 104: 393-411.

Goldman, A. (1992) "Strong and Weak Justification," in A. Goldman, *Liaisons*, Cambridge, MA: M.I.T. Press.

Harmon, G. (1986) *Change In View.* Cambridge, MA: M.I.T. Press.

Henderson, D. (1993) *Interpretation and Explanation in the Human Sciences.* Binghamton: State University of New York Press.

Henderson, D. (1994) "Epistemic Competence," *Philosophical Papers* 23: 139-67.

Henderson, D. (1996) "Simulation Theory vs. Simulation Theory: A Difference without a Difference in Explanation," *Southern Journal of Philosophy*, Spindel Conference Supplement 34: 65-94

Henderson, D. and Horgan, T. (2000a) Iceberg Epistemology," *Philosophy and Phenomenological Research* 61: 497-535.

Henderson, D. and Horgan, T. (2000b) "Simulation and Epistemic Competence," in Hans Kögler and Karsten Steuber (ed.), *Empathy and Agency: The Problem of Understanding in the Social Sciences,* Boulder: Westview Press, pp. 119-143.

Henderson D. and Horgan, T. (2001) "Practicing Safe Epistemology," *Philosophical Studies* 102: 227-58.

Henderson D. and Horgan, T. (forthcoming 1) "The Ins and Outs of Transglobal Reliabilism," in Sanford Goldberg (ed.), *Internalism and Externalism in Semantics and Epistemology,* Oxford University Press.

Henderson D. and Horgan, T. (forthcoming 2) "Epistemic Virtues and Cognitive Dispositions," in Gregor Damschen, Robert Schnepf, and Karsten R. Stueber (eds.), *Debating Dispositions: Issues in Metaphysics, Epistemology, and Philosophy of Mind*, DeGruyter.

Horgan T. and Tienson, J. (1996) *Connectionism and the Philosophy of Psychology*. M.I.T. Press.

Kitcher, P. (1993) *The Advancement of Science*. Oxford: Oxford Univ. Press.

Lyons, J. (1997) "Testimony, Induction, and Fold Psychology," *Australasian Journal of Philosophy* 75: 163-178.

Nisbett, R. and Ross, L. (1980) *Human Inference: Strategies and Shortcomings in Social Judgment.* Englewood Cliffs, N.J.: Prentice-Hall.

Reid, T. (1970) *An Inquiry into the Human Mind*. Chicago: University of Chicago Press.

Sosa, E. (1991) "Reliabilism and Intellectual Virtue," in Sosa (ed.), *Knowledge in Perspective*.

　　　Cambridge: Cambridge University Press, pp. 131-145.

---

[1] I wish to thank Sandy Goldberg for a series of instructive discussions on these issues, and for helpful comments on this paper. I also want to thank Terry Horgan and John Tienson for their contributions and comments on my thought. Finally, I must thank an anonymous referee for this journal, who provided particularly constructive and helpful comments on an earlier draft.

[2] I will here use these terms, 'warranted' and 'objectively justified in believing' as equivalent designations. In so doing, I am not fully conforming to Plantinga's influential use of the term 'warrant'—which, for him, designates that which "epistemizes" true belief (and which must then preclude "Gettierization"). Instead, I use the term as an alternative to talk of objective "justification"—an alternative that may sooth certain misgivings that some feel when reflecting on the way in which talk of justification can bring with it strong associations with fairly rich notions of responsibility, for example (Kitcher 1992).

[3] Of course, conforming to such an understanding might be an important aspect to human epistemic competence. It is arguably central to the ongoing competent refinement of our own cognitive processes. In Henderson and Horgan (forthcoming 2), we argue that certain processes may be epistemically improved by "being under suitable modulatory control" of other processes—and that this kind of joint yoking of processes is commonly required for competent human belief formation, for objective justification. The point seems to us to be related to points Sosa has long advanced.

[4] The case for this is advanced in Henderson and Horgan (2001, forthcoming 1).

[5] In effect, local reliability has come to play a role in discussions of knowledge, as reflected in discussions of false barn regions.

[6] Common standards for epistemic competence are somewhat variable and may best be understood as representing a family of related sets of standards. Some sets of standards may be thought of as characterizing the normal epistemic competence possessed by satisfactory members of a human epistemic community—a kind of satisfactory competence representing the minimum (yet still demanding) set of processes, application of which makes for objective epistemic justification. Other sets of standards may be more demanding in characterizing a kind of

optimizing rather than satisficing normative model of human cognition. See Henderson (1994). The kind of

monitoring here urged as a standard for testimonial belief fixation in normal adult humans has a place in satisficing

as well as optimizing understandings of human cognition.

[7] For discussion of how sensitivity to sample bias contributes to transglobal as well as global reliability, see

Henderson and Horgan (2001) (where transglobal reliability is termed "robustness").

[8] Insofar as the account developed in this paper depends on understanding what processes humans can develop with

training, and which of these contribute to reliability, it goes beyond what is merely conceptually required for

objectively justified belief, to empirical issues regarding what humanly tractible processes satisfy the conditions for

being objectively justified. However, the understanding of these general conditions themselves is generally gotten

by reflection, and is a priori in character.

[9] I must thank an anonymous referee for raising this question.

[10] Here I speak in a commonsensical fashion about the "information possessed" by two agents. Strictly speaking the

formulation would need to be much more complex. The reason is that, as I understand possessed information, it can

be possessed without representation—being "morphological content" in the structure of the trained cognitive system

(Henderson and Horgan, 2000a, forthcoming 2). If there is information in the form of morphological content, then

there is a sense in which agents with significantly different processes in play have and rely on somewhat different

sets of information.

[11] Nor is it to suppose that distinct processes manage these chores: (a) being sensitive to the presence of possessed

information bearing on trustworthiness, and (b) coming to terms with the information possessed. It is merely to note

that the capacity for managing chore (b) when called for requires the capacity for managing chore (a).

[12] See Coady 1992, p. 47, quoted in the text below.

[13] Or processes that are selectively triggered. In any case, it is to demand a background sensitivity, and much of the

paper is devoted to exploring the character of what is required here.

[14] This does not require that, in the satisfied entitlement case, the sensitive processes yield a belief to the effect that

the agent does not possess significant reason to doubt the interlocutor's trustworthiness. Sensitive background

processes could condition episodic processing, befitting the sort of case) without so doing.

[15] I doubt that there are any actual philosophers who really champion AwRI, with its obvious inadequacy. To be fair

to those like Burge who write of default entitlements to accept testimony, one should think in terms of variations on

the basic ARI position that are more demanding than AwRI. AwRI was contrived here to take its departure from the point that *any* position regarding epistemic competence with testimony—even ARI—would need to see simple acceptance yoked with checking processes—and to make vivid the issue of just what these complementary processes must be like.

[16] A qualifiction is needed here, parallel to that in note 23. What is necessary is a reasonable trained up sensitivity— such a trained up sensitivity elaborates background processes in ways we expect and demand of ourselves and others. But, strictly speaking, what is looked for is "acquired sensitivity to real world to real world sensitivity to indicators," for an agent in the real world. In various highly inhospitable epistemically possible worlds—worlds where we would be subject to highly coordinated deception—the training would provide only a sensitivity to apparent indicators of trustworthiness, just as "well-trained" perceptual processes would there produce only sensitivity to much apparent information in a world in which agent's are beset by deceiving demons.

[17] In Burge's classic (1993), memory and testimony are understood as analogous. In memory, the epistemic status of what comes by memory depends on the epistemic quality of antecedent episodes and their traces. In thinking about some problem, I may recall some result. Imagine two variant episodes. In the first episode, I believe that P, having once proved that P, and now recalling that P. In another (indistinguishable as a memory episode), I believe that P, having once erroneously arrive at that belief by a "proof" that was more a matter of wishful thinking than logical competence. The two episodes may be identical, as memory episodes. But the objective justification of the respective beliefs is not identical. My suggestion here is that the epistemic quality of one's reception of testimony (of that episode) can be compromised (and the belief unjustified) if the monitoring or background processes are wimpy by virtue of antecedent negligence. Whether the belief is "in memory" by virtue of past epistemic sins, or allowed to "go through" in testimony because of flawed processes resulting from antecedent sins, it is unjustified.

[18] Fricker terms her own position a kind of "local reductionism." I, however, want no part of the term 'reductionism' in connection with my call for monitoring. I find Fricker's use of the label misleading, as her position allows that monitoring may employ information that includes much information acquired by way of antecedent testimonial belief fixation.

[19] While I argue here for a monitoring requirement on epistemically competent testimonial belief formation by adult human agents, nothing that I say here should be taken to indicate that there is such a requirement on what makes for epistemic competence on the part of very young human children. Such children have not had the advantage of the

long training (in everyday life experiences) with testimony that makes for the sensitivity to trustworthiness that I prescribe here—and any understanding of what makes for epistemic competence in children will necessarily be correspondingly different.

[20] Perhaps it would be preferable to think of simple acceptance as itself realized in a constellation of processes. After all, as described here and characterized in the literature, it involves the interpretive recognition of an interlocutor as having asserted something, and the subsequent formation of a belief "on that person's or source's say so." This view of "simple" acceptance as "already packaged in a bundle" is instructive and helps one appreciate the naturalness of the AM view. For, it bring into focus the idea that testimonial belief-fixation requires an interpretive moment—one in which the agent settles upon an interpretation of the interlocutor's speech act. Interpretation is a complex and sophisticated, if common enough, activity. (This is compatible with the recognition that, in common cases, it proceeds almost automatically. One theme of the present piece is that much epistemically relevant and sophisticated processing occurs in ways that involve the automatic accommodation of much information.) Importantly, interpretation typically involves an understanding of the speaker as having ends and adopting means. Sometimes these are highly generic—as when one takes an interlocutor to be responding to a request for directions with characteristic civility. Other times these can be quite involved—as in understanding an interlocutor to be attempting to close a deal on a new car by raving about the dependability of the 12-cylinder Masaratti engine. The place for interpretive understanding as input into acceptance naturally encourages the view that monitoring is appropriate, for it encourages us to recognize the extent to which information relevant to monitoring for trustworthiness is typically already in play in the interpretive moment itself. Thus, to speak here of a process of simple acceptance is somewhat artificial. However, eliminating this artificial aspect of the present discussion would only strengthen the argument provided here.

[21] One should recognize that much of the information that we have regarding these matters may be difficult to articulate, and yet it may be very usefully and readily deployed. One explanation for this may be that we come to have rather subtle understandings of common belief-structures, social roles, reward-structures and interests, and that we automatically bring such information to the task of understanding the actions, including speech-acts, of those we encounter. These understandings need not be employed in some elaborate discursive argument, but rather may be employed in a kind of dynamic simulation. For some reflection on the information employed automatically in

interpretation, see Henderson (1996) and Henderson and Horgan (2000b). In the epistemology of testimony literature, these considerations are reflected in Lyons (1997).

[22] In truth, this is something of a convenient simplification. When one eventually comes to appreciate how both interpretive understanding of interlocutors and testimonial belief fixation are managed, cognitively, it seems that there just is no discrete cognitive process of "simple acceptance." Interpretation itself, providing the input to the processes that yield the selective acceptance of testimony, always turns partly on issues of motivation and likely errors. As a result, information and sensitivities relevant to what I am calling strong regulation commonly are already in play in the input to the acceptance process. The point is developed in this section of the paper. Simple acceptance is not a discrete process. Rather, it represents a task that any process of testimonial belief formation would ultimately need to perform.

[23] For reasons having to do with the new evil demon problem, and pointing to the epistemic significance of transglobal reliability, these remarks need to be qualified. They hold true of perceptual competence for agents in reasonably hospitable epistemically possible global environments. However, from the epistemic point of view, one can be perceptually competent in a demon-infested global environment, where all one's training has been deceptive, and apparent successes and failures have resulted in an exquisitely sensitive perceptual system that is nevertheless not globally reliable. (See Henderson and Horgan 2001, and forthcoming 1.)

[24] I do not wish to commit to some elaborate settled account of this matter. What I want to insist upon is that perceptual seemings, given or taken, involve what seem to be salient objects standing forth—for example, an animal of familiar kinds situated with respect to the agent and certain familiar things in the agent's immediate environment.

[25] Such processes may commonly work by producing what epistemologists would write of as "reasons" for thinking that interlocutors are (or are not) trustworthy. However, I would not want to limit their operation to the production of reasons—in the classical internalist sense. Such processes might sometimes work by inducing in the agent what one would describe as a vague cautionary misgiving regarding an interlocutor, ones for which one cannot readily articulate "reasons."

[26] The sort of strongly regulated acceptance praised just now for its enhancement of the reliability of testimonial belief formation is monitoring that turns on a sensitivity and understanding of testimonial processes. Such sensitivity itself comes form stretches of life experience, from long training. It is not feasible in early childhood, where normal human dependency on adult testimony does not enable significant choices regarding who to trust. Obviously, it

cannot provide any global vindication of testimony, as some of the information on which it draws will have been acquired beginning in childhood by way of testimony. Rather, it is in the business of refining testimonial belief formation. The refinement that is envisioned would depend on the sensitivities and understandings acquired in the course of one's ongoing training in the ways of others.

Regulated acceptance as described here begins "in the middle," in Jonathan Adler's (1994, p. 269) apt phrase. We and our informant are in the middle of a series of ongoing communications with various sorts of folk (sometimes with the same person). We have understandings of the constraints operating on the informant by virtue of their social role and place within various communities. In even the most passing incident we have some evidence that our interlocutor has engaged in much successful communication (not being institutionalized, presumably, speaking competently in a way that suggest smooth emersion in the community, with all its constraints), and such smooth emersion itself has some (again understood) implications. In these and other ways, we have a great deal of contextual information and understandings that constitute a kind of evidence regarding our interlocutor's trustworthiness.

[27] Interestingly, there are significant limits to the testimonial deception to which an agent in the imagined world is susceptible. In order to deceive the child regarding many things, the handlers must first teach the child the portions of the language necessary to refer to those things. At least with respect to many everyday sorts of things, this requires trustworthy testimony. One must be taught what oranges are, and some things about plant growth and weather before one can be deceptively told that oranges grow well in cold, icy, weather. Such points limit the extent to which testimony can be deceptive, as Burge, Coady, and many others have observed. But, it is not at all clear that they make for as sweeping or powerful an *a priori* guarantee for general trustworthiness as is imagined.

[28] "Explicitly represented" here does not mean consciously represented.