

The Epistemology of Error

Douglas Allchin
Minnesota Center for Philosophy of Science

No question's solved until error's resolved.
—Prospective proverb

1. Introduction. How do scientists know—and justify—that they have erred? The question virtually bristles with paradox. Error seems the very antithesis of knowledge. How could one justify such a "negative" discovery? Oddly perhaps, to *know* that a claim deemed right in one context is wrong requires justification. I focus here on this dimension of the scientific enterprise, the ascertaining of error, and its relation to the general problem of characterizing reliable knowledge.

Error and "negative" knowledge are closely allied. *Negative knowledge* comprises false claims, whose falseness has been justified. It contrasts with positive knowledge, whose claims are also justified, but considered true. *Error* is a false claim interpreted as true and justified. (One can equally imagine the symmetrical case, a true claim interpreted as false.) That is, error occurs when ultimately negative knowledge passes as positive knowledge in some evidential context. In the experimentalist's idiom, error is an artifact interpreted as a signal, or fact. How does one differentiate fact from error, since both seem justified, while only one is true? I am primarily interested here in how false claims can masquerade as (locally) justified.

Some confuse error with uncertainty. For example, scientists refer to *measurement error* and statistically derived *standard error*. Their graphs often include *error bars*. These all denote a range of numerical values, within which the actual value remains uncertain—that is, indefinite or unknown. No error (in the sense just noted) is identified. None may exist. Uncertainty, here, only marks the scope of

potential error associated with a benchmark value. Uncertainty itself (as a form of indeterminacy or ambiguity) is not error. How error and uncertainty do relate is an important theme addressed below (§4).

Claims of error and negative knowledge, I contend, involve the same forms of justification used for positive knowledge (§§2-3 below). Identifying an experimental artifact involves constructing another, alternative fact built in part on the same evidence. To ensure reliability, one must address residual uncertainty by explicitly considering the potential evidence for such alternatives. Even in the absence of anomalies or better explanations, one must probe for error. Resolving fact and error through such further work yields *deep reliability* (§4). Indeed, I claim, knowledge develops by resolving such uncertainties in successive layers. The significance of the epistemic work in ascertaining error and resolving uncertainty guides a program in *error analytics*, which frames many new challenges for philosophers of science (§5).

2. Do We See Mesosomes with a Microscope? Consider the case of bacterial mesosomes (Rasmussen 1993, Culp 1994, Hudson 1999, cited below without dates). Mesosomes are structures detected through electron microscopy, first observed in 1953. Biochemists and bacteriologists searched for their function for over a decade. Microscopists now generally construe them as artifacts, irrelevant products of preparing specimens, not genuine, or "real," structures in bacteria—that is, as error. Why were mesosomes first viewed as fact and later as artifact? This episode helps profile the construction of error and the role of the fact/artifact distinction (Allchin 1997, Latour and Woolgar 1979, Hacking 1984, Galison 1987, Franklin 1996).

Mesosomes were not predicted by any theory. One cannot simply dismiss them as theory-laden observations. They just appeared in electron micrographs. Microscopists certainly knew that they could misinterpret such images. So they calibrated the electron microscope against light microscopes, aware that finer resolution might nonetheless reveal new structures and discrepancies (such as mesosomes). They varied the preparation procedure to see if the phenomenon was robust. It appeared to be. Not that mesosomes were not contingent upon preparation technique. Any new phenomenon needs to be "teased into relief" (as Galison notes). Microscopists gradually developed optimal conditions for revealing mesosomes, exploring the presence of sucrose, glycerol or calcium ions, pre-fix time, temperature, form of cryoprotection, fixative and method of viewing (e.g., thin section v. freeze-fracture) (see Hudson, 306-307). They developed a body of experimental knowledge for producing mesosomes *reliably*. At this level, mesosomes were—and still are—"real." They are wholly reproducible. The conditions for their existence are well characterized: features we do not conventionally associate with error. A stable phenomenon, one might note, need not be meaningful.

In the early 1960s, scientists regarded mesosomes as authentic bacterial structures. Enough so, at least, that biochemists began analyzing their function (Rasmussen, 245-250; Culp, 48). Textbooks featured pictures and diagrams of the prominent mesosomes, noting their role was not yet known. Evidence during this period seemed to warrant acceptance, even if tentative or qualified (a critical historical point missing in Hudson's analysis). How, then, did anyone ever suspect otherwise: that mesosomes were *not* "real"? Why open the mesosome black box?

Reservations emerged in several labs beginning in the late 1960s and early 70s. Nanne Nanninga had been checking the relatively new freeze-etch (now freeze-fracture) technique. Problems

with another structure, the nucleoplasm, arose and were resolved using phase-contrast light microscopy on living cells (Nanninga 1971, 222-23). But for mesosomes, discrepancies between old and new technique persisted. Nanninga (1971) isolated one difference to the use of osmium tetroxide (OsO_4) as a prefixative. But here (she acknowledged) the interpretation of mesosomes was coupled to an evaluation of the methods that produced them. With no independent standard available, one could only withhold judgment (pp. 222-23). M. Silva (1971) echoed her concerns, especially about OsO_4 . Citing inconsistencies from different methods he, too, adopted a new posture of indeterminacy. But he *argued* for this position, appealing to visible differences as evidence. Both investigators began dislodging anchors that established the meaning of the structure in the micrographs. The anomaly of mesosomes emerged, then, during normal science (Kuhn 1972). But the mismatch did not involve theory. Rather, alternate methods generated discordant observations: a *consilience anomaly*.

An anomaly, though, is not a full-fledged error. Interpreting an anomaly requires further experimental work and reasoning. An anomaly signals only the presence of an error. Until it is fully characterized (by isolating it in the procedural-conceptual network), the error is unknown. Confidently accepting mesosomes as artifacts, therefore, involved understanding how they were created. One needed to *explain* mesosomes. One also needed experimental data to *justify* that interpretation.

Such explanatory models of mesosomes emerged in the mid-1970s. Nanninga (1973) hypothesized an enlargement of small membrane features due to "chemical or physical impairment" (pp. 171-74), though without offering any detailed mechanism. She cited relevant information, such as the shape of mesosomes, their placement and the failure of anyone over more than a decade to identify a clear function. M. Higgins' lab based their scheme on the ability (independently demonstrated) of one

fixative, glutaraldehyde, to cross-link proteins, proposing that it caused small, peripheral membrane units to coalesce into one, oversized mesosome (Higgins et al 1976). Margrit Fooke-Acheterrath and her group (1974) again linked large mesosomes to OsO_4 , and showed that chilling could prevent this, but admitted that "the precise mechanism by which the artifacts arise is unknown" (p. 282). All these researchers targeted large mesosomes only. Silva's lab (1976), however, assembled a more comprehensive and thorough model, richly argued with comparisons and controls. First, they showed that use of OsO_4 was correlated temporally with progressive mesosome formation (in both number and size). One could virtually track their development. Further, they proposed a mechanism. OsO_4 damaged membranes, they said. This was observed when they "calibrated" OsO_4 using simple protoplasts (without cell walls). They also measured potassium ion efflux as an indicator of membrane damage, independently of any microscopy. They also considered other chemicals, showing that only those that damaged membranes (not just OsO_4) generated mesosomes. Silva's paper established new standards for interpreting mesosomes. Debate and elaboration followed for at least a decade (Hudson, 301-304), but the ultimate resolution resembled Silva's 1976 sketch: mesosomes are "real," but they are produced only when the bacterial membrane is damaged in preparing cells for electron microscopy. They are not native in the cells. The error, in a sense, was complete.

With Silva's results, the robustness of earlier assessments of mesosomes dissolved. That is, all the previously "diverse," apparently independent experimental methods now collapsed into one common flaw: membrane damage. Different controlled experiments confirmed how this variable was critical. Silva had anchored mesosomes to new benchmarks on a different experimental landscape. They became artifacts: they no longer reflected authentic cell structures. Mesosomes were (are) still fact,

though uninteresting or irrelevant. Of course, these new arguments have their own limits, or qualifications. Mesosomes seem to occur at specific locations. Thus, the pattern of how membranes respond to damage may indicate something (else) about bacterial cell structure.

In summary, do we (did we) see mesosomes with the electron microscope? Using Hacking's (1984) principles, one might say "yes" and "no." Clearly, microscopists identified a "real" phenomenon, separating signal from noise in the spirit of the New Experimentalists. And it was stable, as emphasized by many sociologists (Latour and Woolgar 1979, Pickering 1995). At the same time, Hacking insisted, the mapping from specimen to observed image must be good (p. 320). The mapping ultimately determines the meaning of the image. In this case, mesosomes were *not* good mappings of living cells. Knowing this, however, involved justifying that they are, instead, good mappings of damaged cell membranes. They are now different facts: about how bacterial cells respond to OsO_4 and other treatments. We now *see* mesosomes as error.

3. What Causes Beriberi? The mesosome case represents experimental, or relatively local, error. Error may also be more conceptual, or global. Consider the case of the cause of beriberi (Carpenter 2000). We now view beriberi as a dietary deficiency of thiamine, or Vitamin B₁. But Christiaan Eijkman in 1886 guided his studies using the recently developed germ theory of disease. He began looking for a microbe that caused beriberi. Indeed, the patterns of outbreaks—on ships, in prisons, insane asylums and impoverished neighborhoods—strongly indicated contagion through lack of hygiene. Through a series of accidents Eijkman isolated the cause of a similar disease in chickens to a diet of polished white rice. The polishings, or red coating of the rice, would cure the disease. Eijkman

claimed to have localized the bacterium in the rice, along with an anti-toxin in the coating. To extrapolate his findings to humans, Eijkman and a local official surveyed the incidence of beriberi among the prisons on Java. They identified diets as either polished rice, unpolished rice, or a mixture. The scale of the controlled study was immense: 280,000 prisoners in 100 sites. They also considered and ruled out ("controlled for") other factors that might be microbial vectors: ventilation, age of buildings, permeability of the floors to water, etc. The data dramatically confirmed Eijkman's claims. When institutions later changed their rice diets, the incidence of beriberi decreased. This study capped the work that later earned Eijkman a Nobel Prize.

Though Eijkman's conclusions fit the evidence, they were not necessarily free from error. Other interpretations, outside Eijkman's conceptual horizon, were also possible. Eijkman's successor in Java, Gerrit Grijns, saw the reverse gestalt: namely, something missing rather than something present. He saw the rice coating as containing an essential nutrient. When absent, patients succumbed to beriberi. For him, there was no germ or infection. Contrary explanations, here, each fit the available evidence. Uncertainty resurfaced. Further experimental work was needed. Grijns thus explored the contrast cases. He showed that the nutrient, as a "curative" factor, might be found in other foods, notably the mung bean, *kachang-ijo*. Likewise, non-rice diets of tapioca root or sago might also cause the disease. Grijns *created* anomalies for the bacterial hypothesis, all aligned with his own interpretation. Contextualizing Eijkman's findings in these further results allowed one to locate the error. As no one could isolate any bacterium, despite sustained efforts, others accepted the evidence for a dietary deficiency. Later, biochemists would isolate and characterize more fully the factor now known as thiamine, or Vitamin B₁. Eijkman's error was not obvious. Indeed, his conclusions allowed effective

control of beriberi. Identifying the error involved further epistemic work: first, envisioning the alternative explanation and, then, resolving the different interpretations with appropriate experimental data.

Ascertaining Eijkman's error was coincident with establishing beriberi as a nutrient deficiency instead.

4. Error and the Resolved/Uncertain Distinction. Several observations about these two cases are epistemologically important. They support a new conception of knowledge based on differentiating, or resolving, fact and error from uncertainty:

(1) Justifying error requires epistemological work (Allchin 1997, 2000a).

The histories of mesosomes and beriberi illustrate how one needs evidence to transform a former fact into error. Justification here flows from the same sorts of controlled experiments, calibration, independent background knowledge, robustness, etc., that support any factual claim. Philosophical commentators on the mesosome case disagree sharply about what justification biologists used or was warranted. However, all recognize that error claims need justifying. Paradoxically, perhaps, characterizing an error is a form of establishing knowledge. One must thus beware of *Popper's blindspot*: the notion that falsified claims are non-knowledge and hence discarded. Indeed, collected knowledge of past errors may productively guide research (§5).

(2) Knowledge contrasts with uncertainty, not false claims.

Because knowledge embraces both true and false claims, neither the fact/artifact nor true/false distinctions can differentiate knowledge from non-knowledge. Rather, knowledge (of *both* fact and artifact) contrasts with uncertainty, or indeterminacy—namely, being unable to select

among alternatives. The transition from mesosome-as-fact to mesosome-as-artifact was not like flipping a binary on-off switch. Rather, researchers first retreated to a position of uncertainty. Assessment passed through an "inflection point" of equivocal evidence, epitomized by Silva's and Nanninga's 1971 views. Grijns's challenge was not to show (simpliciter) that beriberi was a nutrient deficiency. Rather, he had to show first how evidence consistent with the bacterial interpretation could also fit a different explanation, too. Then, when combined with further evidence, it indicated something else.

The key epistemological distinction is thus not between true and false, fact and artifact. Instead, it is between empirically unresolved questions, or uncertainty, and resolved questions, where fact and error have been differentiated (with relative degrees of confidence). The primitive state is uncertainty, not being wrong:

positive knowledge, or fact (true)	negative knowledge, or artifact (false)	knowledge of fact <i>and</i> artifact (resolved)
		uncertainty (unresolved)

(a) conventional distinction (b) revised distinction

In the conventional positivist model, knowledge is assembled from particulate observations and held together with logic. Knowledge develops by accretion of new facts: a growth model. In the New Experimentalist vision, knowledge consists of events "carved away" or "removed" from the background—which is then discarded as error (e.g., Galison 1987): a subtraction model. In the image I advocate, knowledge develops through differentiation and increased resolution. Both foreground and background, fact and artifact, positive and negative knowledge, are equally important. What matters is

differentiating claims in successive layers, just as many images are now transferred serially over the internet (Figure 1).




Figure 1. Simple visual model of the resolution of knowledge. Poorly resolved claims (upper left) are differentiated more finely into true and false (here, black and white).

The New Experimentalists and laboratory ethnographers have richly documented this practice, without articulating the epistemology fully. For example, few distinguish clearly between error and uncertainty (e.g., Knorr-Cetina 1999, 276-77). So, while they recognize the importance of separating fact and artifact, they do not recognize uncertainty as the primitive state. Fact and artifact co-emerge from undifferentiated perception. The signal/noise metaphor is closer, so long as one does not confuse noise

(uninterpretable randomness) with negative knowledge. As noted above, mesosomes are not noise. They are a clear signal. But the *wrong* signal. Signals can be fact or error. Latour and Woolgar contrast order and chaos, using a thermodynamic analogy based on entropy (1979, 244-52). The metaphor fits if one underscores the two-fold product. Scientists map negative knowledge as much as fact. The two products are complementary. Indeed, secure factual claims are contingent upon understanding (and differentiating) possible error. Hence, researchers' widespread concerns about addressing sources of error. Understanding the epistemology of error highlights the role of resolving fact and error from uncertainty.

5. Error Probes, Deep Reliability and Error Analytics. Two more conclusions emerge from these two cases (§§2-3). Building on the notion of resolving uncertainty (§4), they lead to a concept of *deep reliability* and a family of epistemic strategies under the banner of *error analytics*:

(3) Reproducibility alone does not establish meaningful fact.

Mesosomes are perfectly reproducible, even today. We know how to produce "good" mesosomes in contrast to "poor" mesosomes. Interpreting them is a distinct process. So, too, can one use an erroneous bacterial model to cause beriberi with rice diet. As Walter Gilbert once cautioned, "you can reproduce artifacts very, very well" (Judson 1981, 170). Mere replication (or lack thereof) differs from interpreting experimental results (contra Collins 1985, 19, 130). One needs some other method to circumvent error.

(4) Verification alone does not guarantee reliable (error-free) conclusions.

Even though results may match precedent or expectations and thereby modestly endorse a method or conceptual model, fact and error may still be unresolved. Demonstrating a model's viability in one context or domain does not warrant universal acceptance. Thus, the locally stable interpretation of mesosomes could dissolve when the scope broadened to include other preparation methods (and biochemistry, too). The meaning of mesosomes changed. Eijkman's controlled experiment, too, had limits. Uncertainty lingered without resolving the results further. Absence of anomalies does not secure the meaning of results. These episodes show how deeper investigation may expose unrecognized error.

Philosophers are all too familiar with the problem of underdetermination of theory by data. These two cases resonate well with Mayo's neo-Popperian remedy (1996, esp. Chap. 6). For Mayo, reliability hinges on a dual process of confirmation *and* ruling out error (pp. 4-7, 184-85, 315). One must actively and aggressively entertain possible error: an *error probe* (pp. 64, 445). For example, one might probe whether hitherto uncontrolled variables are relevant. For Popper, the aim was merely to falsify a theory (or a model or, more broadly interpreted, to invalidate a method). For Mayo, one does not merely selectively jettison certain concepts. Rather, in the spirit of medical diagnostic probes, an effort to falsify can be a constructive tool for arguing (conversely) about reliable fact (p. 183). The strength of one's test (or its *power* in statistical cases) depends on the degree to which one could find error, should it exist. Silva's 1976 studies of mesosomes and Grijns' on diet exemplified well this strategy. They investigated (and effectively demonstrated) that specific variables disregarded earlier were indeed relevant. An *error probe* is the tool for researchers to differentiate fact from error-that-masquerades-as-fact.

Error probes are common in scientific practice, though philosophers perhaps do not yet fully appreciate their role. Scientists' wide concern about error is vividly evident, for example, in their pervasive (and familiar) use of supplemental experimental controls. Checking parallel conditions through a control, they identify whether specific conditions critical for interpreting experimental results are present or absent. Controls of this type help researchers rule out certain variables as irrelevant, or causally unimportant. They exclude alternative conclusions. They deepen confidence in how one interprets experimental results. Thus, the concept of an error probe captures elements of experimental practice not aimed at providing direct positive evidence.

Experimental conclusions that survive rigorous error probes, Mayo says, pass a *severe test* (1996, 64, 178-83). That is, one *argues from error*, in the sense that one considers specific errors experimentally and finds none. Severe tests and arguments from error are not infallible, as my case studies show. There were good reasons, at least initially, for interpreting the images of mesosomes as genuine structures of living bacteria. Eijkman, likewise, seemed warranted in viewing rice as a causative agent, having ruled out (or balanced) other possible factors. (We can easily see *now* how the prison study was not severe, or powerful, enough to rule out interpretations of diet in terms of nutrient deficiency.) One can characterize qualitatively, therefore, the severity of tests. It depends on how one has probed the possible alternative interpretations of results at all levels. A suite of experimental results or controls must invalidate each error. Thus, when characterizing the reliability of their claims, many scientists discuss specific sources of error, not probabilities or scales of belief. The concept of severity resonates with the notion of eliminative, or limited induction, also known as the "Sherlock Holmes

strategy" (Franklin 1986, Earman 1992), Bernard's method of counterproof, and Platt's "strong inference." But the concept of an error probe extends it further.

Reliability is certainly founded partly on verification or agreement between various inductions and experimental results. However, error probes go further by exploring the well known gap between verification and reliability. Error probes yield *deep reliability*, where one has resolved uncertainty between fact and error experimentally.¹ Reliability deepens as one excludes possible error through fuller investigation. Philosophers of science, I contend, will profit from further attention to error and error probes in science. How does the search for deep reliability guides how scientists design experiments, construct arguments and frame criticism?²

The epistemic view based on the resolved/uncertain distinction and adopting deep reliability as a central principle I call *error analytics*. Fully resolving fact and error means probing for error in addition to demonstrative verification. Error analytics complements and generalizes *error statistics*, which examines quantitatively the uncertainty in cases of sampling error and other random phenomena (Mayo 1996). Several maxims may summarize error analytics and the the central tenet of error probes:

Nothing's concluded 'til error's excluded.

Check for flaws before declaring laws.

Uncertainty lasts until probing error is past.

No question's solved unless error's resolved.

An error analytic view invites philosophers to articulate more fully the dynamics of learning from error. What strategies allows researchers to isolate, identify and remedy error, once an anomaly is encountered (e.g., Bechtel and Richardson 1993; Darden 1990; 1991, Chaps. 8, 11, 15)? How do

researchers use knowledge of error? For example, researchers typically develop an informal catalog of past mistakes: an *error repertoire* (Mayo 1996, 5, 18; Allchin 2000c). Where this memory guides avoiding similar mistakes again, standards of proof escalate. Reliability deepens. How should scientists document and communicate errors—or negative results—towards more effective science (Allchin 1999a)? Other prospective projects include characterizing canonical errors, or general *error types* (Mayo 1996, 18, 316, 453; Allchin 2000b). Because errors may be social, as well as experimental, error analytics can potentially unify philosophical and sociological perspectives (Allchin 2000b). I interpret these problems as invigorating a philosophy of science caught between the New Experimentalism and unresolved issues of reliability, and challenged by sociological cases of error. Well construed, an inquiry into the epistemology of error therefore not only highlights the resolved/uncertain distinction and deep reliability, but also may launch a new research programme to guide philosophy of science into a new century.

Acknowledgements. Work was supported in part by NEH #FS-23146 and NSF #SES-0096073.

David Dressler suggested the label 'error analytics' and helped me develop the error theme. I am indebted to Lindley Darden, Peter Lewis and Kent Staley for commenting on earlier drafts.

Notes

¹ The concept of deep reliability echoes Harding's (1991) notion of *strong objectivity*. Both preserve conventional concepts of justification, going beyond them by articulating remedies to specific

weaknesses. The gender and ethnic bias that is central to Harding's account is just one form of error (emerging at the social level) in a more general error analytic perspective (Allchin 2000b).

²On experimental design see Rudge (1998) on Kettlewell's peppered moths, Galison (1987, 64) on the Barnett effect, Franklin (1986, 138-64) on Millikan's oil drops, Mayo (1996, 214-50) on Perrin's Brownian motion; on arguments, see Suppe (1998, 393) on the role of data impeachment in; on criticism, see the role of evidential irony noted both by Alan Gross and Greg Myers.

References

- Allchin, D. (1997), "A Twentieth-century Phlogiston: Constructing Error and Differentiating Domains", *Perspectives on Science* 5, 81-127
- . (1999a), "Negative Results as Positive Knowledge, and Zeroing in on Significant Problems", *Marine Ecology Progress Series* 191, 303-305.
- . (2000a), "To Err is Science", AAAS Meetings, Washington, DC (February 21, 2000).
- . (2000b). "Error Types", [submitted for publication].
- . (2000c). "Error Repertoires", [submitted for publication].
- Carpenter, K.J. (2000), *Beriberi, White Rice and Vitamin B*. Berkeley: University of California Press.
- Collins, H.M. (1985), *Changing Order*, London: Sage Publications.
- Culp, S. (1994), "Defending Robustness: The Bacterial Mesosome as a Test Case", in D. Hull, M. Forbes and R.M. Burian (ed.), *PSA 1994*, Vol. 1. East Lansing, MI: Philosophy of Science Association, pp. 46-57.

- Darden, L. (1990), "Diagnosing and Fixing Faults in Theories," in J. Shrager and P. Langley (eds.), *Computational Models of Scientific Discovery and Theory Formation*, San Mateo, CA: Morgan Kaufmann, pp. 319-46.
- . (1991), *Theory Change in Science*, Oxford: Oxford University Press.
- Earman, J. (1992), *Bayes or Bust*, Cambridge, MA: MIT Press.
- Fooke-Achterrath, M., K.G. Lickfield, V.M. Reusch, Jr., U. Aebi, U. Tschöpe and B. Menge. 1974. "Close-to-Life Preservation of *Staphylococcus aureus* Mesosomes for Transmission Electron Microscopy", *Journal of Ultrastructure Research* 49, 270-85.
- Franklin, A. (1986), *The Neglect of Experiment*, Cambridge: Cambridge University Press.
- Galison, P. (1987), *How Experiments End*. Chicago: University of Chicago Press.
- Hacking, I. (1984), *Representing and Intervening*, Cambridge: Cambridge University Press.
- Harding, S. (1991), *Whose Science? Whose Knowledge?* Ithaca: Cornell University Press.
- Higgins, M.L., H.C. Tsien and L. Daneo-Moore (1976), "Organization of Mesosomes in Fixed and Unfixed Cells", *Journal of Bacteriology* 127, 1519-23.
- Hudson, R.G. (1999), "Mesosomes: a Study in the Nature of Experimental Reasoning", *Philosophy of Science* 66, 289-309.
- Knorr-Cetina, K. (1999), *Epistemic Cultures*, Cambridge, MA: Harvard University Press.
- Latour, B and S. Woolgar. (1979), *Laboratory Life*, Princeton: Princeton University Press.
- Nanninga, N. (1971), "The Mesosome of *Bacillus subtilis* as Affected by Chemical and Physical Fixation", *Journal of Cell Biology* 48, 219-24.

——— (1973). "Freeze-fracturing of Microorganisms. Physical and Chemical Fixation of *Bacillus subtilis*", in E.L. Benedetti and P. Favard (eds), *Freeze-etching Techniques and Applications*. Paris: , pp. 151-79.

Pickering, A. (1995), *The Mangle of Practice*. Chicago: University of Chicago Press.

Rasmussen, N. (1993), "Facts, Artifacts and Mesosomes: Practicing Epistemology with the Electron Microscope", *Studies in the History and Philosophy of Science* 24, 227-65.

Silva, M.T. (1971), "Changes Induced in the Ultrastructure of Cytoplasmic and Intracytoplasmic Membranes of Several Gram-positive Bacteria by Variations in O_2 Fixation", *Journal of Microscopy* 93, 227-32.

Silva, M.T., J.C.F. Sousa, J.J. Polónia, M.A.E. Macedo and Ana M. Parente (1976), "Bacterial Mesosomes: Real Structures or Artifacts?", *Biochimica et Biophysica Acta* 443, 92-105.

Suppe, F. (1998), "The Structure of a Scientific Paper", *Philosophy of Science* 65, 381-405.