Chomsky vis-a-vis the Methodology of Science

Tom Johnston.
Academia.edu. Https://Independent.Academia.edu/ThomasJohnston7
Philosophical Papers. Thomas Johnston – PhilPeople
September, 2016.

Table of Contents

Chomsky vis-a-vis the Methodology of Science....1
Prologomenon....2
Prologomenon Bibliography....3
Keywords....4
Scientific Methodology....4
Rules and Laws: Against Chomsky....4
Rules and Laws: TM0....6
Why Linguistics is Not a Science....9
Elegance and Simplicity vs. Complexity and Constraints....9
What is Required to Be a Science?....13
Intersubjectivity....14
Measurement....15
Law-Law Constraints: Intra-Theoretical....15
Law-Law Constraints: Inter-Theoretical....16
Discovery: Mechanisms and Metaphors....16
The Components of Scientific Metaphors: Bills of Material and Operations Manuals....18
Discovery: Mathematics and Measurement....21
Religion....25
Stories, Laws and Rules....26
As-Yet Unincorporated Material....34

Prologomenon

> Note: this prologomenon was written on October 12, 2121. As indicated above, the paper to which this is the prologomenon was completed in September, 2016.

In the first part of this paper, I review Chomsky's meandering journey from the formalism/mentalism of *Syntactic Structures,* through several methodological positions, to the minimalist theory of his latest work. Infected with mentalism from first to last, each and every position vitiates Chomsky's repeated claims that his theories will provide useful guidance to later theories in such fields as cognitive psychology and cognitive neuroscience. With the guidance of his insights, he claims, psychologists and neuroscientists will be able to avoid costly dead-end lines of research.

This never happened. As I have shown, this never could have happened. (See Johnston, 2018). What has happened, instead, is that current neurolinguistic research (with the arguable exception of the now-dated Lemma Model of Willem Levelt) proceeds without reference to Chomsky. It also wholeheartedly rejects the mentalism of the associated Language of Thought theory of Jerry Fodor. (See Johnston, 2018).

I make this argument in the first part of this paper. I would also like to point out that most of my argument was developed in 1972, when I was a graduate student, lost in the wilds of academic politics, and working on an apartment maintenance crew during the day, and as a tax transcriber for the IRS at night and, later in that year, taking a full-time job as a computer programmer. I know of no other sustained criticisms of Chomsky at that time, and certainly none along the lines I had developed back then.

Although I completed my doctorate in 1978, I never got an academic job, and consequently was never able to publish any of my work in academic journals. So as someone who has had to work beyond the pale of academia all my life, I want to emphasis the originality of my work at that time. But I am now, in my retirement years, using *academia.edu* (and, lately, *Philosophical Papers*) to publicly record the research I began in my graduate student days, and which I have resurrected here in my retirement years. (resurrected in Johnston 2016 and 2018, the currently canonical statement of the semantic theory begun in 1972, and which I call *TM*).

I am proud of this work.

In the second part of this paper, I present my own account of the methodology of science. When I was a graduate student, philosophy of science was dominated by an attempt to describe a methodology common to all the specific sciences, i.e. Hempel's deductive-nomological model. These days, Hempel's emphasis on the methodological unity of

science has been rejected by such "dis-unity" philosophers of science as Ian Hacking, Patrick Suppes and Nancy Cartwright (see Cat, 2021).

I view this change as the swing of a pendulum or, to change the metaphor, a journey from one end point of a continuum to another. As the level of abstraction at which one tries to describe scientific method is raised, the descriptions become increasingly general. Whether or not unity-of-science theories become so general as to be vacuous, is ultimately a subjective judgment. And so I expect that philosophers will eventually become tired of increasingly specific "close to the workbench" descriptions of how scientists work, and begin to turn back to methodological "big pictures", finding in them powerful abstractions rather than empty irrelevancies.

In the second part of this paper, I present my own account of the methodology of science, which I would situate somewhere between the "unity" and "dis-unity" accounts. However, I am not a scientist. My own views about scientific method have three origins:

- my work as a graduate student from 1966 until I passed my comprehensive exams in 1973 (at a different university);
- reading every issue of *Scientific American* from 1972 until nearly 2000, (at which point I continued to read it only sporadically, since I concluded that, around that time, it had evolved from a serious science magazine to a popular science magazine); and
- my three-year immersion in the cognitive neuroscience of language after I retired, based on repeated study of and note-taking for (Banesh & Compton, 2018), (Kemmerer, 2015), several other books and, finally, numerous articles not hidden behind a paywall.

So, as always with my writings: caveat emptor.

Prologomenon Bibliography

Banesh, Marie T. & Compton, Rebecca A. *Cognitive Neuroscience, 5th edition*. Cambridge University Press, 2018.

Cat, Jordi, "The Unity of Science", The Stanford Encyclopedia of Philosophy (Fall 2021 Edition), Edward N. Zalta (ed.), URL = https://plato.stanford.edu/archives/fall2021/entries/scientific-unity/).

Johnston, Tom. "The Co-Ascription of Ordered Lexical Pairs: a Cognitive-Science-Based Semantic Theory of Meaning and Reference: Part 1". Https://Independent.Academia.edu/ThomasJohnston7, August 2018. Also at *Philosophical Papers* – Thomas Johnston – PhilPeople.

Johnston, Tom. "The Co-Ascription of Ordered Lexical Pairs: a Cognitive-Science-Based

Semantic Theory of Meaning and Reference: Part 2".
Https://Independent.Academia.edu/ThomasJohnston7, August 2019. Also at
Philosophical Papers – Thomas Johnston – PhilPeople.

Kemmerer, David. *Cognitive Neuroscience of Language*. Psychology Press, 2015.

Keywords

actors, agents, best theory, big data, bills of material, black-box theorizing, Bloomfield, brains, causes, Chomsky, co-ascription, cognitive neuroscience, cognitive psychology, competence, componential analysis, constraints, corpus linguistics, evolution of language, externalism, falsifiability, generative semantics, generative syntax, holism, innate perceptual gestalts, inter-theoretical constraints, internalism, intersubjectivity, intra-theoretical constraints, Jackendoff, Katz-Postal, laws, learned perceptual gestalts, linguistic behavior, linguistic rules, linguistics, mathematics, mechanistic explanations, Minimalist Program, narrative, Nature, observation, operations manuals, parameters, patterns, perceptual gestalts, performance, plot, Postal, prediction, principles, protolanguage, quantification, Quine, reasons, regularities, restrictiveness, rules, science, semantic primitives, sentences, simplicity, size-complexity continuum, stories, theory of meaning, theory, things, TM.

Scientific Methodology

Rules and Laws: Against Chomsky

Bloomfieldian linguistics was a matter of catalogs. Chomskyean linguistics has been a matter of rules, organized with formalisms equivalent to mathematical logic. But it has never been a matter of laws, and that is why it has never been a science. But besides the pejorative sense of not being a science, and of the laws vs. rules distinction (about which, see below), in what sense is MIT linguistics (as I shall call linguistics done by Chomsky, his followers, and his opponents, i.e. all those who do linguistics "dans le style de Chomsky") not a science?[1]

As just indicated, I maintain that the issue is that of laws vs. rules. A rule *describes* a

1 This is an allusion to a story about the English composer Ralph Vaughan Williams who traveled to France to continue his musical education. On his first day there, his music professor asked him to compose "un petit minuet dans le style de Mozart". The next day, Vaughan Williams returned to England.

regularity; but nothing in the theory of which that rule is a part *explains* that regularity. A law, however, both describes and explains a regularity.

In the early days of MIT linguistics, new rules came fast and furious. And they came easy. Adjustments caused by the introduction of a new rule were merely adjustments to re-partition the descriptive territory to make room for the new rule whenever its applicability seemed to cover not only some newly-discovered syntactic pattern, but also some already-covered patterns that a new rule could more elegantly (or, later on, more restrictively!) account for.

As time went by, attempts, such as Postal's abstract syntax, were made to ascend to higher level generalizations of syntactic regularities, thus consolidating any number of previously different rules. A short time after that, Chomsky abandoned his own principle of elegance as his criterion of theory-choice, and substituted the principle of restrictiveness – the former principle awarding first prize to the theory with the fewest rules and the latter principle awarding it to the theory with the most.

Regardless of Chomsky's methodological vacillations, early and late, the rules of the MIT linguists didn't support one another. How could they? These rules were nothing more than descriptions, as laden as possible with technical terms, of observed syntactic patterns, and given the honorific title of "rule". Most of these rules were transformation rules, but many were not. Some were meta-rules, such as the Katz-Postal rule that transformations do not alter meaning. ("do not"; not "cannot". There is nothing in linguistic theory to explain "can" and "cannot".)

Although Chomsky has long-since jettisoned his performance/competence distinction, I think the explanatory vacuity of linguistic rules is a legacy of that distinction. The main value of that distinction, in fact, was to do what Bloomfield also did, i.e. to dismiss as irrelevant any questions about linguistic behavior that could be investigated by, for example, cognitive psychology and, more recently, also by cognitive neuroscience. These sciences, like all sciences, rely on the quantitative measurement of phenomena, and on

the use of controlled-variable repeatable experiments to confirm hypotheses or disconfirm them or (as Quine would point out) to make adjustments elsewhere in the web of theory and observation that constitutes any scientific theory – or indeed our common-sense theories, too.[2]

The rules of MIT linguistics are neither confirmed by nor disconfirmed by quantifiable data and controlled-variable repeatable experiments. Such data, such experiments, and the evolving web of theory and observation created by such cognitive practices, are irrelevant to MIT linguists. They are theories, hypotheses and data about performance, and competence alone is the game that these linguists play.

Once a theory of competence is well-established, and the relevant cognitive sciences more fully-developed, Chomsky has said, then those sciences will be guided by the linguists' theory of competence. Guided by that theory, those sciences will be able to make rapid progress because the theory of competence will have already eliminated many dead-end lines of research that would otherwise surely have been pursued.

Rules and Laws: TM0

This, however, is just bunk. As I wrote in 1972:

> #1
>
> Chomsky has developed a *rule-oriented theory.* But science works with *laws* which are part of a *regularity-oriented* type of theory. The birth of a science is usually the replacement of rule-talk by regularity-talk. In a choice between rules and regularities, the latter have often proven more useful.
>
> ..... So, we have, Chomsky says, innate rules. And when we talk about consciously formulating rules and following them out, I can understand what is being said. But when we talk about rules, in a context in which we do not consciously utilize them, or judge ourselves and others in their light – this is obscure.
>
> It is always possible to describe any orderly process as rule- governed. This possibility is intimately related to vitalism and teleology. But the possibility of so describing them does not entail that entities are "following rules" (supposing any sense could be given in that

2 An early and beautiful statement of Quinean holism is in the Introduction to his *Methods of Logic* (1959).

context).

..... For we have unlimited instances of processes which science and scientific laws explain, without recourse to the notion of entities following rules. This is the difference between a rule and a law. Laws express regularities that obtain; rules express regularities that are to obtain.

(TM0, [1972-06-16b – Plain], pp. 30-32.)

#2

Another feature of Chomsky's defense of his theories is that he describes phenomena in highly theory-laden terms. Thus, "the child" applies complex rules "on the basis of degenerate evidence".

This is an absurd confusion. The child engages in certain kinds of verbal behavior, exhibiting certain regularities in the production of novel utterances. Chomsky's theory interprets this phenomenon as one in which rules have been applied.

At points, Chomsky does lip service to this distinction. But he continues to describe evidence in transformational-generative terms, thus creating an impression of great empirical support for his theory.

However, described in a non-question-begging way, the evidential support for transformational-generative grammar is less than impressive. The question is this: in face of the evidence concerning language use and acquisition we have, what kind of theory ought we adopt to account for it?

Behavioral and neurophysiological psychology is one possibility; transformational-generative grammar is another. Chomsky correctly points out that the former is not yet developed enough to plausibly account for the normal creative use of language.

But there are crucial methodological reasons for not adopting the latter.

Laws and rules: both are statements of regularities. But laws are statements which are connected to a description of entities, and their properties and relations, which exhibit these regularities. (These entities may be directly observable, or theoretical. The basis for this difference is psychological, though the difference itself is, of course, logical.)

A theory – with its laws, not rules – is a heuristic device which suggests revisions of laws, adjustments, etc. Without it, we could only adjust or refine our laws after the fact. With it, we can anticipate.

Transformational-generative grammar has no such theory because it is a body of rules, not of laws. Thus it cannot do more than adjust, after the fact, to unexpected phenomena. It is not a useful tool for correctly anticipating future experiences.

(TM0, [1972-06-21 – Plain], pp. 36-37.)

#3

When we do not have a law-oriented account of phenomena, a rule-oriented account can be of help. It can give us a handle on the phenomena, and with it, we feel that we know what to do to bring it about again, or to discourage its repeated appearance.

But the rules do not provide an account of the structure of the object or objects engaged in the activity which is the phenomenon to be accounted for. Chomsky acknowledges as much in his black-box theory of theorizing.

Here is what Chomsky does. He describes a class of phenomena, A. He posits a class of phenomena, B, and a set of rules, R, such that a person (or a computer) could generate a member of A by beginning with a member of B and applying some of the rules in R.

There are deep objections to this model of theorizing. A black box approach deprives us of a story about the structure of the objects whose behavior constitutes the phenomena in question, as the story about atoms and molecules explains phenomena like the heat/pressure correlation of a gas, and as the story about elementary particles explains phenomena like the evolutionary patterns of stars. Without such a story, we have no way of distinguishing between alternative sets of rules, each adequate to function as R for some given subset of linguistic utterances.

Because of this black box approach, there is no reason to believe that when the neurophysiological processes which produce the phenomena of linguistic behavior are gradually elucidated, Chomsky's rules will be seen, in retrospect, to have contributed anything to their discovery. Chomsky says his rules are the "abstract structures" of these physiological processes (note 1, Chapter 1, *Aspects of the Theory of Syntax*), and while the absence of any ontological commitment to his theories about linguistic phenomena makes it possible for him to make that claim, no matter what is revealed about the physical processes by scientific investigation, that very absence of commitment makes his claim vacuous. Consider:

Chomsky says that when the neurophysiological processes underlying linguistic behavior are gradually elucidated, his mentalistic, rule-oriented theory will have been incorporated into the body of science; the neurophysiological discoveries will have put meat on the bones of his "abstract structures". This is equivalent to saying that the explanations of disease in terms of modern biology have incorporated talk of demons and evil spirits into science; demons and evil spirits are the abstract structures which modern biology discovered the physical correlates for.

The progress of science is a progress in which explanations by rules are replace by explanations by laws. The laws state regularities relating to physical entities. Rules describe how a phenomenon could be produced by an object capable of following rules. But the phenomenon of creative language use is not the result of a process extended enough in time to be the following of rules. And if unconscious rule following can be instantaneous, then following rules has been divorced from all those conceptual contexts in which it is intelligible.

(footnote: 11/16/2015. But consider playing football. Players follow the rules (mostly), but they aren't thinking about the rules while they are playing the game. The point here should be just that children don't learn language from a grammar book or a dictionary. They don't learn language by learning rules (until later on, when forced to).)

(TM0, [1972-07-19 – PhilNotes72-3], pp. 46-48.)

#4

Chomsky's rules are so elaborate that it is inconceivable that a scientific theory about creative language use will bear even a rough correlation of parts to parts with a Chomskyean set of rules. Consequently, those rules cannot be said, in any reasonable sense, to have been vindicated by the scientific theory, or incorporated by it, or have played a heuristic role in its generation.

(TM0, [1972-07-19 – Green], p.49.)

#5

Linguistic theory is a mentalistic theory. It can retrodict, but not predict. Thus it is not a

scientific theory. What, then, is its use?

I can't think of any very important use for it.

Chomsky's performance-competence sophistry ought not blur the important fact that linguistic theory will not become scientific due to an extension of science, because it is unable to predict, and when neurophysiology begins to account for language use, linguistics will simply be seen as a myth, not a theory of any kind (competence).
(TM0, [1972-08-02a – Green], p.63.)

#6

Chafe says that concepts are neurochemical units (*Meaning and the Structure of Language*, p.75). But this and Chomsky's similar claims, are just whistling in the dark. A witch-doctor could claim that disease-causing demons exist, and are entities specifiable in macro-molecular terms. The attempt at respectability by association, however, fails in both cases. For in both cases, the claim is vacuous.

Under what conditions would such a claim carry weight? Pragmatic ones. Mendel's concept of a gene was a theoretical concept. Genes today are observable.

And the claim that Mendel's genes are the macromolecules observable today carries weight because the Mendelian concept was specific enough to guide the research which eventually led to the modern concept.

I have already pointed out that not only are linguistic concepts not capable of doing that, they actually mislead neurophysiological research because of the theory-laden observation language employed.
(TM0, [1972-08-03b – Green], p.67.)

Why Linguistics is Not a Science

Elegance and Simplicity vs. Complexity and Constraints

So much for the ancient history of MIT linguistics. Ancient history, but history from which we can draw currently relevant methodological conclusions, two of which are the two primary methodological mistakes that Chomsky made.

Recall that, after originally emphasizing simplicity and elegance as ajudicatory mechanisms guiding us to the "best theory", and finding that Postal's abstract syntax bested his own theories on that very count, Chomsky switched to complexity and constraints ("restrictiveness") as the better criterion.

Take, first of all, simplicity and elegance, Chomsky's earlier of these two antagonistic

methodological positions. The problem with simplicity and elegance is that Nature does *not* select for the simplest and most elegant solutions. It selects for solutions to whatever problems arise, using material already at hand.[3] Consequently, the same genetic components appear over and over again in the descent of species, e.g. using the same genes for insect legs, amphibian flippers, birds' wings, and human arms. This cross-speciational re-use indeed, is the core message of evo-devo.

In doing so, some of Nature's most impressive constructions resemble Rube Goldberg (Heath Robinson) contraptions more than the highly-engineered products of current technology. The extremely inelegant positioning of optic nerves in the human eye is an example.

The greatest of Nature's solutions is human language, a product of the human brain. But by placing linguistic behavior beyond the pale of the competence/performance distinction, linguistics has (until recently) developed without reference to those neural mechanisms and processes by which language is produced. Linguists who then argue that their elegant and simple theories will eventually guide neurolinguistic research overlook the simple point that Nature works with the genetic material at hand, on problems that must be solved soon enough to permit reproduction to continue the genome.

Mutations, of course, do introduce genetic novelty. But successful mutations are a needle in the haystack of the sum total of mutations within a genome, most of which vanish before the first resultant phenotype can make its appearance, and most of the rest shortly after that.

3 "If it ain't broke, don't fix it!" is not a recipe for progress in the engineering of human artifacts. But Nature is quite happy to proceed on that basis, and also on the basis of a companion principle, "If it is broke, patch it!".

Take, secondly, complexity and constraints. The problem with complexity and constraints in linguistics is that both are artificial. As for complexity, the Principles of P&P are whatever principles linguists make up. As for constraints, the Parameters of P&P are whatever constraints linguists make up. Speaking of the rules of MIT linguistics in general, no set of those rules constitute constraints in any meaningful sense because nothing constrains what constraints linguists can make up as they go along.

Furthermore, nothing in later work in the MIT tradition seems to involve any constraints at all. As focus shifted from generative syntax to generative semantics to lexical semantics, for example, it was no longer transformational rules (or lexical derivations) which were made up as the theorist went along. It was semantic primitives.

But what constrains the choice of semantic primitives, beyond the imaginations of linguists?[4]

Perhaps, the response to all this might be, that a different game is being played by linguists, the real game being "Give me my assumptions and let's see where they lead us, what range of linguistic intuitions they seem to elegantly (or restrictively!) account for."[5]

That game, it seems to me, is, in fact, the one Chomsky has played from the beginning. He then changed the rule book and emphasized constraints and restrictions, after abstract syntax beat him at his own "elegantly accounts for" game.

4 See, for example, Jackendoff's ontological categories (1983. *Semantics and Cognition*), and Souvran's use of "the notions of 'difference', 'continuity', detachment', 'membership', 'identity', 'direction', 'value' and others (as) crucial, formative-meaning elements that organize the words in domains and determine the differences and the nuance between their members". (2014. Souvran, *Relational Semantics and the Anatomy of Abstraction*, p.20).

5 Assumptions about transformations, about abstract syntax, but also assumptions about lexical primitives. These latter assumptions range from the components of Katz and Fodor's componential analysis (and where did those components come from?) to present-day relational semantics. They range from assumptions made by both the interpretative semanticists and the generative semanticists, by, for example, both Jackendoff and Lakoff (so methodological blame here accrues to both camps).

But then and now, everything about linguistic theory is reverse-engineering without the engineering, all retrodiction and no prediction. There is no hypothesis-testing in linguistics, no crucial experiments in which generally-agreed on relevant variables are held fixed while one variable is allowed free play. And although I agree with Quine's holistic point that it is a theory as a whole which is being tested in every crucial experiment, in practice scientists generally agree on what part of a theory is verified (or falsified) by an experiment.

None of this happens in linguistics.

The recent interest in corpus linguistics, and especially in compiling data about the co-occurrence of ascriptions of lexical items, is a welcome move towards replacing Chomsky's isolating rationalism with a more balanced and more science-like research program of formulating theories with testable hypotheses, testing those hypotheses, revising the theories as seems best at each revision occasion, and iterating that process.

But in recent work in corpus linguistics, although we have gotten back to actual data about language at work, where is the theory? Is it that two lexical items which are frequently ascribed together are related by meaning? If so, this simply shows an ignorance of conclusions already reached and settled in the philosophy of language. The famous example is Quine's "creature with a heart" and "creature with a kidney" (and, of course, Frege's morning star and evening star).

The problem for corpus linguistics is that lexical items can be co-ascribed to a statistically significant degree *either* because they describe things which go together in the world, *or* because patterns of co-ascription have become so widespread that violating them will be taken as evidence that the person doesn't understand the meaning of one or

the other of those predicates. The former is empirically-based co-ascription; the latter is semantics-based.

Extension and intension both influence lexical co-occurrence, and they are very different. The trick is to explain the difference and then to weave the differences together.

Corpus linguistics provides good data (though, as anyone familiar with Zellig Harris' work will appreciate, that data has to be put into some semblance of "canonical form" before it can be cranked through a statistics machine). But if what we're looking for is meaning, where is the theory of meaning which this data may be relevant to?[6]

TM is such a theory of meaning. It explains the difference between co-ascriptional occurrence which reflects semantic commitments, and co-ascriptional occurrence which does not. It is an internalist (meaning in the head) theory of meaning which explains externalist intuitions such as that the meaning of sentences has a lot to do with whether or not the world is as those sentences say it is.

What is Required to Be a Science?

But I am getting ahead of myself. The previous section is negative, explaining why linguistics is not a science. But what, then, is required for it to be a science? My topic, in this section, is a brief sketch of what it is to be a science, interspersed with comments on how linguistics fails to meet these criteria.

A note: this is not just "We're scientific and you're not!" name-calling, the kind of thing with which the history of the Linguistic Wars is replete. Pragmatically speaking, it is an appeal to linguists to consider what kind of theories they produce and what the value of

6 See (2014. Lenci. "Will Distributional Semantics Ever Become Semantic?"), and my commentary (2016. "Notes on Lenci on Distributional Semantics").

those theories are, and a suggestion that those theories (especially the ones focused on syntax) are, in fact, no more than catalogs of features of a language (English, almost exclusively), different from the Bloomfieldian catalogs that Chomsky derided only by being dressed in fancier clothes.[7]

What distinguishes science from non-science, as I see it, are six factors:

1. The intersubjectivity of observational data;
2. the precision of observational data, achieved by quantification;
3. the web of constraints which laws impose on one another, and on new laws;
4. the use of mechanistic explanations, i.e. explanations in terms of things and the impact of those things on one another, in theory discovery and theory evolution;
5. following where the mathematics leads as another guide to theory discovery and theory evolution; and
6. prediction (and retrodiction) of precise (quantified) and intersubjectively verifiable observations that follow with, of course, an implicit certeris paribus clause implemented, to the extent possible, by holding fixed variables that could, in a theory-consistent way, explain observed outcomes, while allowing those processes to occur whose changes are recorded in the values of the theory-relevant variables not held fixed.

Intersubjectivity

Intersubjectivity of observation is a matter of perceptual gestalts which are conditioned to names of objects of interest by learning, and standardized by agreement and correction in

7 Another note. This account of what it is to be a science is the product of my own personal reflection over several decades. It is not based on experience as a working scientist. Nor is it based on research in the philosophy of science, either late-20th century work, or more recent work such as that which focuses more on differences among different sciences than on formulating a grand unificatory theory such as the hypothetico-deductive method.

the presence of the phenomena with others in the same community of interest. Intersubjective agreement on labels attached to the things we perceive (and to the discrimination of those things from their backgrounds) is prerequisite, and where it is not given, can be learned. Given that, we can move on.

Measurement

Precision of observation is the result of quantifying observations. Quantifying began with counting, and it is only discrete things that can be counted. By attaching scientific instruments to events, we can count many things that otherwise we could not count. Also, with the development of the calculus, we can measure continuously varying processes, and can determine the value of any point in a continuously varying field.

Law-Law Constraints: Intra-Theoretical

Laws in science express intersubjectively verifiable perceived patterns, more or less directly tied to observations (quantifiable patterns of gestalt-organized sensory input), but none in isolation from the system it belongs to, which includes other laws.

Rules, as used in linguistics beginning with Chomsky, are also expressions of (generally) agreed-upon perceived patterns, originally focused on syntax, but subsequently extended to semantics as well. But none of these rules are quantified, since linguistics for many decades ignored linguistic behavior. And, even more significantly, linguistic rules do not provide constraints on other linguistic rules, because they can be made up willy-nilly, as they have been since Chomsky. And the reason they, and not scientific laws, can be made up willy-nilly is that scientific laws, unlike linguistic rules, are part of a theory which has experimentally testable and measurable (and so falsifiable) predictions of what will be observed when the experiment is conducted.[8]

8 Let's say, for now, that my use of the term "MIT linguistics" refers to linguistics from *Syntactic Structures*, to *Aspects*, to x-bar theory, to Principles and Parameters, to Government and Binding, and not to exclude the Generative Semantics movement and the generativist/interpretavist debates, but to exclude the Minimalist Program and Chomsky's later speculations on the evolution of language.

Law-Law Constraints: Inter-Theoretical

I refer here to what I called, in TM0, the size-complexity continuum. This is the continuum which defines the subject matter of the physical sciences, from quantum physics to macro-physics to relativity, from quantum physics to organic chemistry to biology, and from quantum physics to inorganic chemistry to the molecular engineering of the materials sciences. Along this continuum, there is a requirement that laws for a higher-level science must be consistent with laws for the next-lower-level science. For example, chemical bonding theory must be (and is) consistent with the quantum-mechanical interpretation of electrons as waves (fundamentally as waves, as theorized in quantum field theory) whose observed localizations are interpreted as particles (the electrons themselves) with a varying probability of occurring anywhere within the atomic orbitals to which Schrodinger's wave function and Planck's constant confines them.

Discovery: Mechanisms and Metaphors

An important feature found in nearly all scientific theories is explanation of things in terms of their parts, and of the actions of things in terms of the interaction of those parts in response to external stimuli caused by contact with other things. For example:

- Natural environments have parts, and ecology is the study of the interaction of those parts and the effects those interactions have on the states of those parts and on the state of the ecological system as a whole.

- Within an ecology, organisms exist. Organisms have parts, and biology (and medicine) is the study of the interactions of those parts (organs), and of groups of those parts (the hormonal system, the gastro-intestinal system, the muscular-skeletal system, the nervous system, etc.) and the effects those interactions have

on the states of those systems and on the state of the organism as a whole.

- An organism is made up of organs, which are made up of both structural and dynamic organic molecules, and whose processes are carried out by means of creating and consuming organic molecules. Organic molecules have parts, and organic chemistry is the study of the interactions of those molecules and the effects those interactions have on the states of those molecules and on the creation or dissolution of specific molecules as they come into contact with other organic molecules, breaking and forming valence-shell bonds. At the interface of organic chemistry and cell biology is the study of the interaction of these organic chemistry objects and processes with the cells that make up the organs that make up the organisms that we all are.

- An organic molecule is made up of atoms consisting of electrons, protons and neutrons, in three-dimensional shapes, and which by means of electron repulsion and attraction among molecules can be aligned with other molecules in a "lock and key" configuration, at which point the closely-fitting atoms in the lock and key area can modify the congruent molecules by means of creating and breaking electron bonds, thus reshaping the participating molecules, breaking some of them up, and creating new ones out of the participants, in the process consuming and releasing energy.

- Atomic structure (the elements) is based on the number of protons and neutrons in an atom. Chemical behavior is based on the attraction and repulsion of electrons in the outer electron orbits of atoms. Chemistry rests on atomic / quantum physics, and the laws of the former are required to be consonant with the laws of the latter. In other words, the laws of chemistry are restricted and constrained by the laws of physics.

- Electrons are negatively-charged particles, with a wave function that describes the probability of an observation finding an electron at some given point in its orbital. Electrons emit and absorb photons as they change orbitals, thus losing and gaining energy.

These levels of explanation of physical nature constrain one another. No laws in organic chemistry, for example those involving the transformations of acids into bases and vice versa, will be considered if that law violates quantum theory, for example by requiring/assuming a different configuration of electrons in an atom's electron shells than is determined by the Aufbau principle.

The Components of Scientific Metaphors: Bills of Material and Operations Manuals

These parts usually have some structure. For example molecules are three-dimensional structures made of up atoms, which are made up of electrons, protons, neutrons, and so on. In manufacturing, structure is represented by a Bill of Materials, which is a tree-structure in which under the root node (which represents the item whose structure is being depicted), all the child nodes are parts which make up their immediate parent nodes.

If we are interested in more than structure, if we are interested in process and change, then a Bill of Materials will not suffice. In that case, we need one other thing: an Operations Manual. The manual describes how something works – both the internal changes that happen to it, and the external changes that its workings bring about. We will need to itemize the inputs to those processes, as well as the internal states of the objects involved in them. We may need to itemize the parameters of the environmental conditions under which the processes can proceed. Given all this, observed processes should be repeatable processes. One scientist should be able to reproduce the results of another scientist's experiments.

There are a few flies in this ointment. One of them is sensitivity to initial conditions, popularly referred to as the "butterfly effect". Small variations in input, or initial states, or working conditions, can defeat the attempt to reproduce a process exactly.

Another fly in the ointment is the inability of some experiments to measure a process without disturbing the process being measured. A famous example is the two-slit experiment with light in which photons will pass through slits and form a wave pattern if not observed with an instrument which detects which slit each photon passes through, but if observed will not form that wave pattern.

Another example is the mechanism of all microscopes, optical or electron. Both types operate by the reflection of a particle (a photon, an electron) from the object being observed. But this reflection alters the state of the object being observed. At the atomic level, it alters the location and/or vector (direction and momentum) of the atom. This is the well-known Heisenberg indeterminacy in which increasing accuracy in measuring the location of a particle results in decreasing accuracy in measuring its vector, and vice versa – this being because to observe a particle, a photon or electron must collide with it.

And we need to itemize the outputs of the process, including the change of state of the participants.

In many cases, statistical reasoning on a large number of as-near-as-possible repeated processes will be needed to establish the correlations between variations in these features of the process and variations in the outputs produced.

What I have had in mind here are the physical sciences. When I mentioned the manufacture of human artifacts, I was thinking of the physical objects involved in it –

objects (in manufacturing) classified by where they appear in the process, those classifications being raw materials, works-in-progress, and finished goods. We have raw materials, works-in-progress, and finished goods in science, too. We call them inputs, the intermediate states of objects participating in non-yet completed processes, and the final states of the participants in and products of those processes.

The social sciences have a more difficult problem to deal with than the physical sciences because, by definition, human beings are *agents* in the processes studied by the social sciences.[9] And except when human beings are regimented into being organic cogs in

9 I suggest that we consider molecules and cells as *actors* in the processes of life, but not as *agents*. This distinction between actors and agents is one I intend to develop, in some later piece of writing. For now, the distinction I have in mind can be characterized in a few different ways. The first is this: *actors take part in processes; agents initiate processes*.

For example, in statements like "Language speaks man", for all its appearance of profundity, what is going on is to confuse actors (speech acts) with agents (originators of speech acts). But we do not normally attribute agency to speech acts, and if we are to do so much violence to the semantics of the concept of agency as to do so, a lot of semantic work will have to be done to integrate the consequences of this disruption into the semantics we all share. Lacking that work, statements like "Language speaks man" are just nonsense or, at best, meaningful at the level at which poetry is meaningful (which is, basically, the language of non-cashed-in metaphors).

And the confusion is even deeper than that, for, it seems to me, Heidegger has also reified language, so that he is attributing agency not to speech acts as such, but to an agency-capable *thing* – language. In which case, he at least accepts that agents are things. But agents are things that initiate actions, as when a person, the agent, decides to stop off at the grocery store on his way home from work. Treating language as a thing is an ontological stretch, but perhaps defensible. But treating language as an agent isn't just a semantic stretch; it's a semantic break.

When talking about human beings, to treat someone as an agent is to be willing, for purposes of understanding, to attribute his actions to will and decision, and to hold him responsible for them. Suppose he just kissed his wife, took out the trash, and came back and had a drink. He had a *reason* for doing each of those things. But what neurophysically caused him to do each of those things is irrelevant to understanding *why* he did each of them. And to ask why he had those reasons, i.e. what beliefs leading to actions they represented, and what led him to acquire those beliefs, is out of place (expect in such specialized contexts as psychotherapy sessions). Reasons are explanatorily originating links in the behavior of things capable of having beliefs and acting on them. Reasons answer the question "Why?", and deflect that question when it is directed to them.

But I have proposed a more general sense of agency than specifically human agency. In that more general sense, we don't attribute will and decision to non-human agents, but we do something similar. For purposes of understanding what we are presently concerned with, let us say that a sequence of events, involving different physical objects, starts with an action of an agent. Suppose that agent is a rock that has been thrown at me. To treat that rock as an agent is not to impute intentionality to it. Rather, it is to treat it as the start of a chain of events we want to understand. The *person* who threw that rock, let us suppose, is an originating object in a chain of reason-based explanation and, as such, an agent. But the *rock itself* is an originating object in a chain of physical-causation-based explanation and,

industrial processes, the initial states of those human beings make any mechanistic account of their behavior extremely difficult to obtain. Increased sophistication of brain-scanning technology is opening up vast new ranges of insight into the working of the human brain. But in general, for the social sciences, reliable data relies on statistics applied to human behavior under controlled conditions, ideally, but usually under less than perfectly controlled ones.

Discovery: Mathematics and Measurement

Quantum physics is that branch of science where, at least for the time being, our stories break down. These stories, such as the story reaching from the Bohr model of the atom, to organic chemistry, to cell biology, are broad, deep and coherent. As such, they guide

---

as such, is also an agent. What the person and the rock have in common is this: both are the originating links in a chain of events which we consider adequately explained as long as we can follow the links.

So a second way to distinguish actors from agents is this: *actors are objects at work within a system; agents are objects which initiate processes from outside that system*. Where, for purposes of explanation, we draw the inside/outside line, is what makes the distinction between agents and actors, not vice versa.

In this more general way of making the distinction, the neural sciences can make use of it. Sensory stimuli originate with things external to the organism (or, at least, to its brain) and those things, as such, can be thought of, treated, conceptually accessed, as agents, as origin links in an explanatory chain. As those stimuli affect neurons, they initiate processes in brains. Those neurons, then, are actors, and via transmission along their plastic circuitry, they affect other actors. This cascade of effects are neural processes.

Perhaps this also provides a point of view on the distinction between *reasons and causes*. On this view, the difference between reasons and causes is not ontological, or even epistemological, but "perspectival".

To accept an explanation of a person's behavior in terms of *reasons* is to treat as an explanation a statement of the purported beliefs the person had which led to that behavior. Reasons are unexplained explanatory originations of an intention-capable agent's behavior. To treat a person's behavior as reasonable is to tacitly accept the explanatory rule that it is extraneous to the understanding we are trying to gain to search for an explanation of why the person had those reasons (i.e. beliefs in their role of motivating behavior), or how he came to have them.

To accept an explanation of a person's behavior in terms of *causes*, on the other hand, is to include, in the scope of our quest for understanding, things and processes, either outside or inside the person, which can themselves be treated as resulting from the causal influence of other objects and behaviors on them. So a *cause* is a link in an explanatory chain for which it is legitimate to search for antecedents, if one chooses to do so. A *reason* is a link in an explanatory chain for which it is not legitimate to search for antecedents.

research interests, suggesting questions to ask which lead to further research. That is, the story records and interrelates patterns which have already been discovered, and suggests that new patterns might be found by experimentation and research in certain directions.

If the work is done and the suggested patterns are found, the story has been extended. If not, the tension between the failure of discovery of the anticipated patterns, and the existing story, reverberates through at least localized chapters of that story, and stimulates the work needed to reconcile the experimental results with the story. So the development of this particular and magnificent story is not one of smooth progress; what we thought we already knew will often be modified as reconciliation of unanticipated consequences progresses.

But with quantum physics, its story of waves and particles, of mass particles and force-exchanging particles, breaks down. So when physics graduate students (and some experienced physicists) get too concerned with the conundrums of wave/particle duality, quantum entanglement, and energy in the vacuum of "empty" space, they are often told to "shut up and calculate!" For the quantifiable precision with which quantum physics predicts outcomes of experiments is unrivaled anywhere in science. The standard theory works, even though we don't understand it, i.e. even though we cannot fully incorporate these well-known conundrums into the basic mechanistic story of sub-atomic mass particles affecting one another by the exchange of force particles, or into the so-far distinct story of curved space and variable rates of time.

Yet quantum physics continues to make progress. The mass-conferring Higgs boson was hypothesized because the story, as of a few decades ago, seemed to require it. And more recently, as a conclusion from sophisticated statistics compiling millions of observations made in the LHC, it has apparently been found.

But having just written this, I have discovered how easy it is to speak misleadingly, even

to oneself. Because the search for new patterns in quantum-physical phenomena is driven, not or not primarily, by its background story, but by the pure mathematics of quantum physics. Einstein's equations produce infinities when applied to quantum phenomena, yet relativity is well-established by repeated and increasingly precise experimentation and measurement. So in generating infinities as the solutions of equations which otherwise describe observable processes, the mathematics points research in the direction of eliminating those infinities, of unifying relativity and quantum physics. And in raising the question of the differences in mass of the different sub-atomic particles and their families – quantifiable, measurable differences – the mathematics combines with the basic elements of the physics story, i.e. mass particles (electrons, protons, neutrons) affecting one another by exchanging force particles (photons, gravitons), to point the way to the need for a mass-conferring particle.

So in its search for new patterns, and its need to explain existing ones, quantum physics has deviated from a reliance on stories, and from the need to be able to explain the results of research and experimentation in terms of its story. Quantum physics follows the equations in its search for new patterns and explanations of existing ones. Chemistry and biology do, of course, produce equations and quantified measurements, such as dissociation constants of different molecules and the role of the chemical processes which those constants describe in the regulation of chemical processes. But these mathematically expressed patterns can all be understood as patterns of the actors and processes which are already part of the story of chemistry and biology.

Thus the quest to discover new patterns, and to produce an integrated explanation of the obviously (or presumably) integrated reality which manifests them, can be guided in either of two ways. In one way, this progress in understanding reality is guided by knowledge of an ever-deepening and ever-widening story of things and their interactions. In the other way, it is guided by a description of those patterns which abstracts from their physical embodiments and is formulated in purely mathematical descriptions of those patterns.

These two ways of discovering new patterns, and correcting our understanding of mistakenly perceived ones, have something in common. And this gets back to the rules vs. laws distinction.

Any description of a pattern, in a science guided primarily by a story of things and their interactions, must fit into the overall story. We can't explain an unexpected outcome of a chemical reaction, for example, by saying that the hydrogen bonds were stronger than expected. That hydrogen bonds are weak bonds, weaker than ionic or covalent bonds, is a well-established fact within the story of chemistry. Some other explanation, one consistent with the story, or at least with the most well-established parts of the story (such as hydrogen bonding strength) must be found. In addition, the mathematically-expressed patterns in chemistry and biology cannot be ignored. Explanations must be consistent with both the story of chemistry and with the equations of chemistry which have been already amply confirmed.

In a science in which the story falls behind the mathematics, as it currently does in quantum physics, explanations of new patterns must be consistent with the mathematics. No equation in physics stands alone, and together, the web of equations has proven to be as powerful a guide to the discovery of new patterns as the original mechanistic story of physics prior to the quantum.

It's scary to let go of our stories, but so far it has proven safe to have done so. Nonetheless, I hope that eventually, the story of physics will catch up to the equations of quantum physics, and become capable of providing a coherent description which includes those phenomena which are currently quantum-physical conundrums which don't yet fit into the story.

And this is what the two guiding forces of the scientific investigation of patterns have in common. In the explanations of science, everything relates together; nothing stands alone – the mirror of a reality which we assume, and must assume, itself relates together, a reality in which nothing stands alone. Stories are involved in these explanations, to the degree they can be, because stories are the only way human beings, until the scientific revolution, have been able to make sense of things. And making sense of things is just understanding how they "hang together", in Wilfrid Sellars' sense of the term. Mathematics is the new element, added to stories, that science introduced. Until quantum physics, the only role of mathematics was to provide a quantitatively precise expression of patterns which, in their precision, allowed otherwise unnoticed additional patterns, and apparent deviations from patterns, to be noticed. But with quantum physics, we can see that the mathematics alone can stand on its own as a description of patterns, and can also function as an indicator of where the search to extend that web of pattern descriptions should continue on.

Both stories and equations, again, have this in common. Both express patterns, and point the way to the discovery of new patterns. And both are more than collections of individual patterns. The patterns of science hang together. The belief that all patterns descriptive of reality also hang together is an article of faith with all human beings, but particularly with scientists.

Religion

Religion has functioned, in human history, as a framework for honoring the letter but not the law of this fundamental article of faith by attributing patterns we cannot otherwise explain to an intention-capable agent, i.e. to God. Religion relates patterns by claiming that they are all a result of the will of an omnipotent and inscrutable being. Nothing constrains the attribution of these patterns to God's will and purpose, or prevents the attribution of other, even apparently irreconcilable, patterns to the same source.

Perhaps the patterns of reality *are* ultimately attributable to an omnipotent and inscrutable agent. But without the web of constraints of patterns on one another, we have no guidance. We don't know where to look next. We don't know how to eliminate smallpox, or send machines to Mars, or blow ourselves up. We only know how to hope and pray.

And this is what's wrong with linguistics. It is a collection of God's rules, without God, a collection of descriptions of discrete linguistic patterns, and even of generalizations covering any number of such descriptions. But each rule stands alone in a crowd of rules. Chomsky to the contrary, these rules do not constitute a web of constraints. They constitute a list of arcane descriptions of isolated, armchair-discovered, linguistic phenomena.

Stories, Laws and Rules

In 1972, I developed the concept of *stories* as a way to distinguish between the *laws* of one kind of theory, and the *rules* of another kind of theory. The first kind is a scientific theory. The second is the kind of theory found in the Chomskyean linguistics of the time (and afterwards, at least until Chomsky reached the Minimalist Program).

These first kind of stories are stories about physical things and how they affect one another. In the physical sciences, these things are located on what I called a *size-complexity continuum* (SCC). On the SCC, we have a continuum of sizes, from sub-atomic particles to galaxies and, in fact, extending even further at each end of that range. And we have a continuum of complexity, from photons to brains.

Physical things affect one another by exerting force on one another. Ultimately (as far as our current theories go), these forces are exerted by the movement of force particles from one sub-atomic particle to another – the particles corresponding to the four physical forces of gravity, electromagnetism, the strong nuclear force and the weak nuclear force.

The different physical sciences study different ranges along the SCC. Let's say that each one has its own story. Each story must be internally coherent, and verifiable by repeatable experimentation and observation. But each story must also fit into the collection of stories that make up the SCC narrative. In doing so, each story reinforces other stories, for example (i) the Bohr atom story being reinforced by the wave function probabilistic interpretation of electron obitals, (ii) the entirety of chemistry being reinforced by the Bohr atom story, (iii) the transmission of signals across a network of neurons being reinforced by the story of three-dimensional molecules affecting one another's valencies where areas of those molecules are brought close enough for those electrical forces to have a valency effect, electrical forces consistent with the Bohr model of the atom and the Schrodinger wave function, and so on.

Laws and rules are both statements of patterns – patterns in structures and patterns in behaviors. Ultimately, these patterns are all variations on a single theme, that theme being the patterns described in mathematics. Looked at the other way around, the equations of mathematical systems describe patterns at a level which entirely abstracts from their physical realization in any particular system of objects and forces along the SCC.

Laws are statements of patterns exhibited by objects and forces located somewhere on the SCC. The behavior of organisms exhibit SCC patterns. The use of language by human beings is a mode of behavior of those organisms.

Laws in the physical sciences are not, of course, laws in a *prescriptive* sense. They aren't constraints which the physical world *must* obey. Rather, they are laws in a *descriptive* sense. They describe patterns manifested by objects and by the interactions of objects with one another.

In doing so, they tell a story about those objects. If a set of objects and interactions can be identified precisely enough to make a perfect, or at least extremely good, distinction

between objects within the set and those not within it (in other words, can be described in such a way as to make that distinction), and between interactions within that set and not within it, and if structural and/or behavioral patterns can be found within that set and quantified, then those descriptions become scientific laws.

It is more difficult, it should be noted, to pick out force-mediated behavior from the background in which it occurs than it is to pick out physical objects from the background in which they occur. The reason is that things go on – interactions occur – in a world in which all sorts of other things are going on at the same time, some of them in the same place or close to the same place. These other interactions and forces may influence the outcome of any given interaction. This is why science conducts experiments in which whatever concomitant variables that current theory suggests might influence an interaction being studied are isolated and controlled. And if that can't be done, an attempt is made to observe as many possible interactions in the physical world as possible, with as wide a range of postulated extraneous variables as possible, and then to compare the results and, with statistical analyses of large-enough numbers, to extract a common pattern of behavior from a common set of participants.

A set of laws for any such set of things/interactions, that has been repeatedly confirmed by experimentation in different circumstances which current theory says should not affect the outcome of the experiment, is accepted as an established set of laws, a description of how that particular bit of Nature really works.

Perhaps we should distinguish between the plot element in a story and the narrative element. A *plot* for a story is a pattern followed by the participants in the story. The *narrative* of a story is focused on particular participants, particular interactions they have, and how those interactions change those participants.

Scientific laws are elements in plots which make up a Borgesian library of science. Each

time a specific narrative (an experiment) shows actors of the kind described in the plot, engaging in the ways that the plot says they should be engaging in, then the plot is confirmed, and belief in its laws is strengthened. If someone believes the plot is only make-believe, she will attempt to construct a narrative in which the actors are as the plot requires them to be, the interactions are of the kind the plot is concerned with, but the outcome is not what the plot says it should be. There is a prince, and a princess he awakens with a kiss, but they do not live happily ever after.

Chomskyean linguistics, in its many protean forms, purports to tell a story. But the patterns in its story are described by rules, not by laws. There are no actors and behaviors; they were ruled out by the competence/performance distinction. Even after Chomsky dropped the distinction, there are still no actors and behaviors. There are, at least since the Minimalist Program, structures and states of a black-box object which generates acceptable sentences. But that Chomskyean speaker-behind-a-veil-of-ignorance is located nowhere on the SCC. In particular, although that speaker would be identified as a human being, nothing about him as an individual or a member of our species, or a member of various human social groups, is relevant to any of Chomsky's hypotheses. Nothing about neurological structures and processes (other than that they changed, at some point in the past of our species, in response to a genetic mutation) is relevant to any of Chomsky's hypotheses. Nothing in paleontological studies of humanoid and non-humanoid communicative behavior is counted for or against his theories. Nothing in studies of human linguistic behavior did, or does (or ever will) count for or against his theories (competence/performance was the firewall here, keeping scientific evidence at bay, and still is even though the distinction is no longer a club that Chomskyeans use to beat up others with).

So, Chomsky's claim, in (1972. Chomsky. *Studies on Semantics in Generative Grammar*) – right on the heels, to be noted, of Postal's apostical article "The Best Theory" (1972) – that his (Chomsky's) MIT linguistics program should continue to (as if it ever did) subject itself to "restrictiveness", i.e. to some kind of objective constraints, is not only vacuous,

being honored entirely in the breach; as a statement about what Chomskyeans were doing at that time, and had been doing, it is false.

Here's what I said about this issue back in 1972:

> When we do not have a law-oriented account of phenomena, a rule-oriented account can be of help. It can give us a handle on the phenomena, and with it, we feel that we know what to do to bring it about again, or to discourage its repeated appearance.
>
> But the rules do not provide an account of the structure of the object or objects engaged in the activity which is the phenomenon to be accounted for. Chomsky acknowledges as much in his black-box theory of theorizing.
>
> Here is what Chomsky does. He describes a class of phenomena, A. He posits a class of phenomena, B, and a set of rules, R, such that a person (or a computer) could generate a member of A by beginning with a member of B and applying some of the rules in R.
>
> There are deep objections to this model of theorizing. A black box approach deprives us of a story about the structure of the objects whose behavior constitutes the phenomena in question, as the story about atoms and molecules explains phenomena like the heat/pressure correlation of a gas, and as the story about elementary particles explains phenomena like the evolutionary patterns of stars. Without such a story, we have no way of distinguishing between alternative sets of rules, each adequate to function as R for some given subset of linguistic utterances.
>
> Because of this black box approach, there is no reason to believe that when the neurophysiological processes which produce the phenomena of linguistic behavior are gradually elucidated, Chomsky's rules will be seen, in retrospect, to have contributed anything to their discovery. Chomsky says his rules are the "abstract structures" of these physiological processes (note 1, Chapter 1, *Aspects of the Theory of Syntax*), and while the absence of any ontological commitment to his theories about linguistic phenomena makes it possible for him to make that claim, no matter what is revealed about the physical processes by scientific investigation, that very absence of commitment makes his claim vacuous.
>
> *.....The progress of science is a progress in which explanations by rules are replace by explanations by laws. The laws state regularities relating to physical entities. Rules describe how a phenomenon could be produced by an object capable of following rules. But the phenomenon of creative language use is not the result of a process extended enough in time to be the following of rules. And if unconscious rule following can be instantaneous, then following rules has been divorced from all those conceptual contexts in which it is intelligible.*
>
> (TM0, [1972-07-19 – PhilNotes72-3-pdf], pp. 46-48.)[10]

10 For other discussions in TM0 of stories, rules and laws, see also the entries: ([1972-06-22a – Plain], [1972-06-29 – Plain], [1972-06-29 – PhilNotes72-2-pdf], [1972-07-11 – PhilNotes72-3-pdf], [1972-07-19 – PhilNotes72-3-pdf], [1972-07-19 – Green], [1972-07-21 – Green])

Proceeding on:

- MIT linguistics is still not a science. It still postulates rules, not laws.
- Well, ok, it's not a physical science, Chomskyeans might say. But we linguists never thought it was (although the neurophysiology of language is certainly a field of physical science). It's a social science, being a science about human beings and their linguistic behavior. (Well, about human beings; but in MIT linguistics, studying behavior crosses the competence/performance divide, which is still there even though its name is no longer uttered.)
- So does MIT linguistics qualify as even a social science? Well, a hallmark of social science is reliance on statistics. Since we can't provide a physical science account of specific instances of human behavior, we go "Big Data" and look for patterns in large numbers. Sometimes we can express these patterns quantificationally, even if we lack a story to explain them. This is exactly what Newton did in his *Principia*; he described gravity in mathematical form, but he fingoed no hypotheses. He produced a precise description of a pattern. But he did not produce a law (in the sense I have defined it, and in the sense that nearly all pattern descriptions called scientific laws conform to).
- Corpus lexical semantics is, by definition, based on populations of observations to which statistical techniques can be applied. It is a nascent social science, in search of a theory. What I want to do is provide a broader and deeper theory of meaning than is currently available, and that will account for important patterns found in those corpuses.

So back to stories and plots, in particular those which are part of the grand interdisciplinary narrative of science as a whole. Suppose all this is true. Well then, so what?

One result is simply to become more aware of what it takes to be a science – not in the honorific sense, but in the pragmatic cure cancer and explain the origin of the universe

sense. We shall then be less content with a rule-based theory masquerading as anything worthy of serious attention.

The other reason is that I want to offer an hypothesis about the neurological foundation of stories, stories scientific and otherwise, stories in the sense of accounts of what's been happening, or what might happen, or who's been involved, where it happened and when, how something was done, etc. This to include plain ordinary stories, little and big, told in passing or written down to later become part of the conceptual background of ordinary life, or grander stories destined to become part of the Western Humanities/Science Canon. This account is that all our stories are selections of neurologically-embodied semantic threads from the currently active (and ever-changing) subset of the neurologically-embodied lexical network in our brains.[11] This "account" is a meta-story, using the story of the biological sciences to explain humanity's use of stories.

Aside from this being an interesting hypothesis that might lead to valuable research in the future, it bears on the language of thought issue as well. In this meta-story of mine, reasoning, in ordinary thought and language, is following a narrative thread through a neurally-embodied semantic network. The rules of formal systems of logic are mathematical-level abstractions from these "following from" patterns in human linguistic behavior. These patterns are, in turn, dispositions to accept movement from a given set of neurologically-embodied lexical units to a new set of units, along a connectivity relationship whose end points are the analytic and the synthetic, and which is assembled by syntax into statements.

If this is true, we can begin to build bridges over the chasm from a logical study of thought to a physical study of thought.

11 A quite-detailed map of which is provided by fMRI studies, and presented in (2016. Alexander G. Huth, Wendy A. de Heer, Thomas L. Griffiths, Frédéric E. Theunissen & Jack L. Gallant. "Natural Speech Reveals the Semantic Maps That Tile Human Cerebral Cortex". *Nature*, vol. 532 (April 28, 2016)).

As-Yet Unincorporated Material

As to mechanistic explanation, in physics it is a matter of forces and fields, the mass particles that are relatively stable configurations of condensed waves in fields, and the force particles which maintain or disrupt those stable configurations. In chemistry, it is a matter of electromagnetic forces of atoms which govern their assemblage into three-dimensional configurations – molecules. In neurobiology, it is a matter of the electrochemical flow of ions among neurons. In cell biology more generally, it is a matter of chemical messages, either targeted or broadcast, either intracellular (second messenger molecules) or intercellular messages among cell components such as the nucleus, mitochondria, the cell wall, etc.

In linguistics, rules are not based on mechanisms, let alone mechanisms on what I called, in TM0, the size-complexity continuum. Therefore, there are no constraints on what rules can be made up to explain a new class of linguistic phenomena, or to restrict the applicability of other rules so that apparent exceptions to them are eliminated by ruling them outside the scope of the original rule. Linguistics, that is, is not a science, and is not constrained by any science.

Linguistics needs a mechanistic model, as do all sciences. The things it posits must be, if not eliminable, then reducible to things in the continuum of physical science. The patterns it observes must be measurable, if only statistically.

Chomsky, especially in (1972. *Studies on Semantics in Generative Grammar.*), emphasized the importance of constraints and the restrictions they impose for any scientific theory. Since his constraints can be made-up, with no inter-theoretical constraints (and, see below, no intra-theoretical constraints other than logical consistency), his constraints, which continue to be used in current linguistic theory (as "rules") are no constraints at all.

Linguistic theories (especially in syntactical studies) are not based on evidence. Instead of intersubjectively verifiable observational reports (whose intersubjectivity is based on sensory input being mediated by both innate and learned perceptual gestalts and on conditioned lexical responses to those gestalts), linguistic theories incorporate judgments of grammaticality, judgments which notoriously differ, in numerous cases, among trained linguists, among men in the street, and between both. Rules in linguistic theories are constrained by neither mechanisms nor mathematics, that is, by nothing at all except the intuitions and aesthetic sensibilities and proclivities of the linguists making up those rules.

Perhaps, the reply would be, linguistic theory is a different kind of scientific theory. And we have to be careful here to keep the discussion from degenerating into "It is a science!" vs. "It isn't a science!" name-calling. So consider this question: what explanatory capabilities do rule-based grammars have that earlier Bloomfieldian grammars did not have – supposed capabilities with which Chomsky trashed Bloomfield? More to the point, in what sense are rule-based grammars explanations at all? Unconstrained generalizations can and usually do retrodict; and lacking quantification, they are often used in the informal prediction/anticipation that guides social behavior day to day. But for precise, falsifiable predictions, where is linguistic theory? It's not in that game. It's not a science.