

# Universals and Induction by David Kelley

1988

TOC Product Description:

**The Foundations of Knowledge** by David Kelley.

The "Foundations of Knowledge" is a 5-lecture series which remains to this day the most thorough and profound explication and analysis of Objectivist epistemology. (7 audiotapes)

[www.objectivismstore.com/p-78-foundations-of-knowledge.aspx](http://www.objectivismstore.com/p-78-foundations-of-knowledge.aspx)

## ***Lecture 4: Universals and Induction***

“Unlike other species, man can generalize and extract universal truths applicable to things he does not perceive. The two cognitive processes fundamental to this capacity are concept formation and induction. In Part I of this talk, Dr. Kelley applies Ayn Rand's theory of concepts – set forth in her *Introduction to Objectivist Epistemology* – to the important philosophical problem of universals. In Part II he concerns himself with the question of justifying inductive inference – specifically "inductive generalization."... An enlightening explication of Rand's theory on concept formation, and why Dr. Kelley regards it as a major contribution to philosophy – a solution to the problem of universals.”

---

## Table of Contents:

<b>Part I: Problem of Universals.....</b>	<b>2</b>
<b>Concept Formation and Universals....</b>	<b>2</b>
<b>Classical Theories of Concepts.....</b>	<b>5</b>
Moderate Realism .....	5
Nominalism.....	6
<b>Rand’s Theory of Concepts.....</b>	<b>9</b>
Concept Formation.....	10
<b>Question Period on Concepts .....</b>	<b>11</b>
<b>Part II: Problem of Induction.....</b>	<b>15</b>
<b>Validating the Law of Causality .....</b>	<b>15</b>
<b>The Method of Induction.....</b>	<b>20</b>
<b>Question Period .....</b>	<b>24</b>

In 1988, Dr Kelley gave a series of lectures titled The Foundation of Knowledge at the Portland Institute Conference.

Lecture 4 in this series was titled *Universals and Induction*.

Lecture 4 was transcribed from audio tapes by Roger Bissell on a request by Dr. Milo Schield.

This transcription has not been updated, edited or reviewed by Dr. David Kelley.

Dr. Kelley has given his approval for the public dissemination of this document.

Headings and items in braces { } were added.

Last updated: 10/28/05

## Part I: Problem of Universals

Thank you very much, and let me add my own welcome to Bill's. It's not very often that I get to talk about epistemology on a Saturday, and at this particular hour on a Saturday, it's all I can do to get my dog to listen to me. Let me make sure first of all everyone has one of the handouts for the lecture. They were available downstairs. And if you don't, there are some additional ones up front.

You know, it's often said that man differs from the animals in his fundamental relationship to his environment. Other species have to adapt themselves to their environment; man can adapt his environment to himself. That's a kind of truism, but it's true also of man's form of cognition. The other species are limited cognitively to what is given to their senses. They have no ability to integrate what they perceive in such a way as to extract information that is not directly given to them. Their perceptual world is organized in terms of concrete, specific relationships and mostly spatial relationships among objects. They cannot reorganize their perceptual world in such a way as to discover new relationships. But man has that ability, and the essence of that ability is his capacity to generalize, to notice common properties and ways of acting, to classify things accordingly, and to extract universal principles and truths that are applicable to things that we do not perceive.

Now, there are two cognitive processes that are fundamental to this capacity, and those are concept-formation and induction. I'm going to talk about concept-formation today and induction tomorrow. We'll see that these two capacities are very tightly related.

### Concept Formation and Universals

Now, my text for today's lecture on concept-formation is Ayn Rand's work, *Introduction to Objectivist Epistemology*. Ayn Rand always said that her theory of concepts was one of her major contributions to philosophy. I certainly agree with that assessment, and my goal today is to explain why I think that assessment is accurate.

Now, when I first read this book, it all seemed pretty clear and straight-forward to me, or to be honest, most of it seemed clear and straight-forward, after the second or third reading. But not having studied any philosophy at the time, I really didn't understand its significance. And I found that – this has been a common experience with many Objectivists – it's clear what the answers in the book are, but it's not clear what the questions were. In fact, her theory is addressed to a profound and profoundly important problem in

philosophy, the Problem of Universals. Philosophers have been wrestling with this problem since the time of Plato and Aristotle, and I believe she essentially solved it. That's the significance of her theory, and that's what I'm going to try to convey today.

So, what I'm going to do first is describe in some detail what the Problem of Universals is, and second to describe the main historical theories of universals, so that we can see what problems philosophers have gotten themselves into. And having done that, having established that context, I think we can see very clearly why Ayn Rand's work is so significant.

So, let us begin by talking about the conceptual level of awareness, in order to formulate the problem. Let's distinguish the conceptual level from the perceptual. Sense perception gives us the awareness of particular concrete objects and events in our environment. We're aware of their specific attributes and relationships. We see their specific actions. We can remember some of them and recall them later, recognize them when we see them again. All of these are perceptual-level abilities. But human beings also have the ability to classify things together into groups, the categories, to form concepts for general types of things and for the abstract properties that they share.

The obvious reflection of this is language. Most of the words...individual words in a language reflect concepts in one way or another, they stand for general classes or kinds of things or attributes, actions, relationships, etc. Our language embodies a vast and elaborate conceptual framework, a system of classification. All right? And we apply this system of classification automatically to a great...to all the things that we perceive, without even having to think about it. When you take an apple from the lunch table, there's a very good chance, I would say the probability is on the order of 100%, that you've never seen that exact apple before, that particular apple, and yet you recognize it immediately as an apple, and you apply to it everything you know about apples. You know, on the one hand, that it's edible and, on the other hand, that you can't have a romantic relationship with it.

Now, that was an example of deductive reasoning. All apples are edible. This is an apple. Therefore, this is edible. All of our deductive reasoning revolves around our concepts; it's an application of concepts to particulars. But concepts are also involved in the other broad category of human reasoning, namely, induction. A child burns himself by putting his hand in a candle flame and draws a generalization: flames burn, so he won't do it again. A crucial step in the inference is his classification of the thing that burned him as a flame. If instead, he had classified it as a yellow object, then he would have drawn a different generalization, that yellow objects burn; and he would have begun acting very

peculiarly in the presence of buttered popcorn, blondes, gold coins, which would be very sad.

The problem arises when we ask what concepts stand for in reality. In perception, if we ask what is it in reality that we're aware of, well, the answer is pretty simple: it's this. The podium that we all see exists in reality apart from us as a single unit. Now, there are philosophical questions about the validity of our awareness of it, and I don't want to diminish the nature of the problem – I wrote a whole book about it – but really, this is a problem only philosophers have. It's perfectly clear what the object of perception is. But this is not so clear in regard to concepts. A concept like "dog" stands for an indefinitely large collection of animals of a certain type. Now, these objects, the Dog Kingdom, so to speak, never stand before us in their totality, nor do individual dogs wear signs indicating their category membership. We have to assemble them into a category by some process of integration.

And the same is true of a concept like "red," which identifies a particular attribute. These concepts also stand for any number of things, any number of red things, but here we have an additional problem. When I describe a car as "red," I'm not describing the whole of its identity, only one aspect, its color. But a thing's identity does not appear before us in perception all divided up into individual aspects. A thing's identity is not like a bag of marbles: the color marble, the shape marble, etc., which we can just pull out one at a time and examine individually. Some process of selective attention is required in order to individuate individual aspects, and when we think about that, the same is true really for a concept for a type or kind of thing like "dog," even though the concept like "dog" stands for the whole animal and not just one of its attributes. Still, not every aspect of its identity is relevant to its membership in the category "dog." So, once again, there is some active process of selective attention involved.

So, the question is: how do we do this? How do we take things apart into their component attributes, and how do we put them together again into categories? And what justification do we have for doing it? What basis do we have in reality? Is there something in the objects themselves that requires us to do this, or at least justifies us in doing it? Perhaps there is no basis. The doctrine of Conceptual Relativism says that concepts are subjective devices. Different individuals or different cultures divide up the world and classify things differently, and no one system is objectively better than any other system, all right? Now, this is an issue of profound importance in philosophy as in science and, indeed, in all of human knowledge. All of our knowledge depends on our reasoning and, as we just saw, all of our reasoning depends on our concepts. Any doubt about the objectivity of concepts will cast doubt on the objectivity of knowledge in general.

Now, before we turn to the various theories that deal with...that address the problem I've just outlined, I want to take a deeper look at the problem itself. There are two important features of concepts that define the problem, and I want to talk a little bit about the nature of those two features. I'm referring to the **universality** of concepts and their **abstractness**. A concept is universal, whereas its instances or members are particular; and a concept is abstract, whereas its instances or members are concrete, specific, determinate. Let's make sure we understand what these terms mean, and let's begin with universality.

Imagine a child learning a language, his first language. Even when he gets the idea that the sounds he hears his parents making have the remarkable property of referring to objects – and this is a major achievement on his part – even when he gets that idea, he still has to learn that they don't refer to objects in the same way. Now, if the child points to an animal wandering around the house and asks...indicates that he wants to know what it's called, he might get one of two answers. He might be told "that's Spot," or he might be told "that's a dog." He hasn't understood these answers until he understands that "Spot" is a name for this particular animal, whereas "dog" is a name not only for this animal but for any other animal of the same type in the neighborhood, any other animal of that same type that he will ever encounter, any other animal of that type regardless of whether he encounters it at all. So, the first and obvious point about a concept, then, is that it is *universal*: it refers indifferently to a number of different individual, particular objects.

Now, these objects constitute a set. And if it were just a matter of putting things together into sets, there would be no particular problem. We can group things together to some extent even at the perceptual level. We can attend to the forest as well as to the trees. But unlike a perceptual grouping, a concept is open-ended. It includes not just the particular instances that one happens to have encountered, but all objects of a certain type. So, we can't learn a concept by memorizing its instances. We have to learn some kind of rule or something like a rule for determining what is included and what is excluded from the concept.

Moreover, we are not usually concerned with those objects as a set in their totality. When I describe something as a dog, I don't care how many other dogs there are. I certainly don't know how many other dogs there are. I'm thinking about this dog, and what I'm doing is identifying its nature. But the point is that what I'm describing also seems to be present any number of other animals. In describing any of these other animals as a dog, I'm attributing to them the same nature, the same characteristics. The question is: what is this nature? What is this thing that many different objects

have in common? And that brings us to the next point, the abstractness of concepts.

The idea that the instances of a concept have something in common suggests that they are identical in some way. But they're not identical. When we form a concept, putting objects together into groups, the objects are similar, but they're not identical. The concept "red," for example, includes a range of perceptibly different shades of red. The concept "dog" includes animals that differ in every respect. Obviously they differ in attributes that are not relevant to their biological identity, attributes like whether they're owned by somebody or not, whether they know how to fetch a stick. But dogs differ also in those attributes that are relevant to their biological category. All dogs have hair, but dogs differ in the color and shape of their hair; all dogs are social animals, but as you know if you ever had a paper route or done door-to-door work, dogs differ a good deal in their sociability; dogs have presumably a common DNA structure which biologists will someday tell us about, but even so no two dogs except identical twins have exactly the same DNA. So, in general, the instances of a concept are not identical, they are merely similar.

A concept for an attribute designates a particular dimension of similarity, and a concept of a type of entity like dogs identifies a number of dimensions of similarity, hair, shape, etc. Nevertheless, despite the fact that the instances are merely similar, we treat them as if they were identical. When I say that dogs bark, I'm treating all dogs as if they were identical. I'm ignoring all the differences among them. I'm treating them as interchangeable units. And when I say that they all bark, using the concept "bark," I'm ignoring all the differences among the sounds that they make. In order to do that, however, I obviously have to ignore the differences among these animals, the differences that are perfectly obvious to me at the perceptual level.

I believe this is one of the reasons, this fact of ignoring differences, is one of the reasons why people sometimes resent being classified. People sometimes say "don't pigeonhole me," when they're described under some category such as "conservative" or "housewife" or "engineer" or "athlete" or whatever. The classification presupposes or seems to convey that the person is identical with anyone else in that category and ignores the unique character of the person.

But be that as it may, the fact that we can integrate individual objects into a single category and treat them as identical allows us to present in the form of a single proposition – to use the same example, "dogs bark" – a fact that would otherwise require many different propositions, "poodles yap," "boxers growl," "Dobermans snarl," etc. And even that is stated conceptually. I used the narrower concepts, "poodle," "Doberman," etc. Ultimately, to avoid concepts, I would

have to say "this dog makes this noise," "that dog makes that noise," and so on for all 47 billion or however many dogs there are in the world. By treating instances as identical, the concept allows us a tremendous cognitive economy. It allows us to reduce a mass of information to a single unit.

Now, let's pause here to make sure we understand the difference between universality and abstractness. And I think the clearest way to summarize the difference is, in terms of a philosophical distinction, between what's called **numerical difference** and **qualitative difference**. If I take two coins of the same denomination, these happen to be American quarters, these are numerically different, they're two objects, but they're qualitatively identical, at least as far as perception goes. [I can't...] They are exact duplicates of each other, whereas a quarter and a nickel or any other coin would be qualitatively different as well as numerically different. {The American quarters are not qualitatively identical. They vary in certain perceivable aspects, e.g., how clean they are, what year they were minted, etc., otherwise they couldn't be united into a concept!}

Now, a concept is universal because it includes an open-ended range of numerically different objects, numerically distinct objects. A concept is abstract because it includes an open-ended range of qualitatively distinct objects, objects that are merely similar but not qualitatively identical. {This implies that "American quarter" is universal but not abstract and is thus not a concept, while "coin" is both universal and abstract. Does Kelley mean to imply this?} A concept is universal because it unites a numerical diversity. A concept is abstract because it unites a qualitative diversity. {Again, Kelley seems to be saying that "American quarter" is not a concept, because it is only universal and not abstract, because what it unites is not a qualitative diversity.}

And I think you can see that abstractness is, therefore, the more fundamental of the two features. The reason that we can integrate numerically different objects into a single mental structure is that we can ignore their qualitative differences, that is, we can unite their different qualitative things. So, a concept is universal in virtue of being abstract. So, the Problem of Universals, in my opinion, really should have been called the Problem of Abstractions.

Now, this Problem of Abstraction has two components, a metaphysical and an epistemological aspect. Metaphysically, as we saw, the instances of a concept are similar but they're not identical. But what is similarity? Ayn Rand describes it, and this is a fairly standard description, as: "Similarity is the relationship between things that have the same attribute in different degree."

But what is this attribute that things can have in different degree? Two different shades of red possess the same attribute of redness, but what kind of thing is this common color, redness? Suppose I said, “I just bought a car, and it’s red,” and you said, “Oh, what shade?” and I said, “No shade. No particular shade. It’s just red. It’s bare red.” Or, suppose I said, “I have a dog,” and you said, “What breed?” and I said, “No breed,” and you said, “Oh, you mean it’s mixed,” and I said, “No, it’s neither purebred nor mixed bred. It’s just a dog, standard-issue dog.” Well, clearly this is impossible. An attribute such as “red” or a general nature such as “dog” has to exist in some particular, concrete, determinate form. So, then, the question is: what’s the metaphysical status of the abstract attribute or kind, essence, such as “dog”? And what relationship does it bear to the particular concrete forms in which it exists?

The epistemological problem is to explain how, by what process we ignore the differences among things in order to form an abstract concept. If you show me two dogs and tell me to ignore all their differences, all the respects in which they differ, well, I’m going to have to ignore their shape, their color, their size, the length of hair, their way of behaving, their gender, their degree of docility, their preferences in food – in short, I’m going to have to ignore all of this, and what’s left? Nothing. I’ve thrown out every attribute that they have. If I have to ignore all their characteristics in order to form a concept, then the concept will be completely empty, all right? So, a theory of concepts has to solve this epistemological problem, explaining how it is that we can ignore the differences without ignoring the object as a whole.

### Classical Theories of Concepts

All right, then, let’s turn now to the classical theories that tried to solve these problems; and I’m going to start with the theory of Realism. Realists hold that our concepts mirror the structure of attributes in reality. If our concepts are universal and abstract, that’s because something out there in the world is universal or, at least, abstract. Concepts refer to real universals, Redness, Humanity. These are always spelled with a capital...always capitalized to indicate that they have a special metaphysical status.

Now, there’s something...there’s an obvious problem here in saying that there are these real universals. If Humanity, for example, exists as something outside the mind, it must be some definite particular thing. How, then, can it be shared in common by all the particular human beings? If Humanity is what’s essential to you and me and all the other human beings, and it’s a single thing that we both have, how can you and I and all the other individuals be different

individuals? Perhaps we all share a slice of the universal. This is a theory that Plato once proposed or entertained. You might have rationality; I might have animality. But no, that doesn’t work. All of us are fully human; we’ve got the whole shot, all right? So, how can a single entity be shared, literally shared, by many different particular entities?

Well, this problem led Realists in two different directions. Extreme Realists, of which Plato is the best example, said that the universal Humanity is not in any of the particular human beings. It’s a separate entity, existing in a realm beyond particulars, the world he called the World of Forms. Particulars that have the same attribute, such as red, or belong to the same type, such as human beings, do so because they relate to the same universal, Redness or Humanity, existing in a different world. The universal is shared by all instances in the sense that they all bear the same relationship to it, in the way that many mirrors, arranged around a candle would all have an image of the candle in it. They would all reflect the same one entity.

### Moderate Realism

Now, the other approach, taken by philosophers who were unwilling to posit a supernatural realm, is called Moderate Realism, and the best example of this is Aristotle, at least Aristotle as normally interpreted. Moderate Realists deny that universals exist in some separate realm. Indeed, Moderate Realists deny that anything outside the mind is literally universal at all. But they say that particulars do contain as intrinsic elements in their identity abstract essences and attributes. My essence, my humanity, is distinct from yours, but because our essences consist of abstract properties, they are qualitatively identical. My Humanity and your Humanity are like the two quarters that I held up before, numerically distinct but exactly alike in the two of us. So, in that sense, it’s the same attribute that we share.

To get a better feel for the metaphysics here, let me give you an analogy. Imagine several cars of the same make as they leave the factory. Now, to all appearances, these cars are identical. But once they’re purchased and driven for a while, they each acquire their little individuating elements. One gets souped up, one gets fuzzy dice hanging off the mirror, another one gets a bumper sticker, and so forth. They get customized. Well, for the Realist, when two objects of the same type come off the metaphysical production line, so to speak, they are identical, they’re just red or just human, but then they get particularized to a specific, concrete, determinate form of the attribute, when they’re embodied in matter, when they come to exist as material objects. So, if we take two shades of red, they differ in their specific shade. To perception, they are merely similar, not identical, and yet the

similarity is a product of two factors. We can analyze it into two factors. There's an element of identity, in that the presence of the same abstract feature Redness in the two of them; and there's an element of difference, an individuating element that makes the two shades different. The point to stress here is that the abstract feature exists as an intrinsic element in the identity of the color. And in the same way, my Humanity, the abstract essence, exists as an intrinsic feature of me, and the same with you.

Now, Realists explain the cognitive process of concept-formation in terms of these abstract properties and essences. The properties exist out there in reality as intrinsic components of objects, so concept-formation is like perception. We are directly aware of what is there before us. In perception, it's the whole object. In concept-formation, it's the abstract property or the essence. Now, some preparatory of sifting and comparing similar objects may be necessary before we can isolate the abstract feature that they have in common. We may have to focus our attention in a certain way in order to perceive, in order to be in a position to see the abstract essence or property, just as at the perceptual level you may have to train your attention a certain way in order to see a fracture, a bone fracture in an X-ray. But once you've trained your attention and focused your attention in the right way, then you just see it. And the same, says the Realist, is true of concepts; once you have trained your attention in a certain way, you just open your mind, and there it is, the abstract property that the concept refers to.

So, in this way, as you can see, the Realist can explain why the instances of a concept are treated as identical. It's because they *are* identical; they all share the same identical property, and we are directly aware of that property. Unfortunately, this is exactly what makes Realism untenable. The Realist doctrine of intrinsically abstract attributes runs into all sorts of problems as soon as you start exploring it in depth.

One classical problem is the Problem of Borderline Cases. On the color spectrum, red shades into yellow. Now, if what is common to all shades of red is the feature Redness, and what's common to all shades of yellow is the abstract feature Yellowness, then at some definite specific point, Redness must end and Yellowness must begin. But it seems arbitrary to stipulate any one point where that happens. And if you try to solve the problem by saying, "Well, there's orange in between," so there's red, orange, and yellow, but then we have the problem about red and orange. I mean, you're not going to solve it by getting narrower concepts.

Second problem. How many distinct attributes does an attribute have? Just staying with the realm of color, is an object's color a single abstract attribute, Redness? Or, is it the three attributes that psychologists

study, Hue, Brightness, and Saturation – does it have three abstract attributes? Or, does it have the many elements of color that a painter can recognize and reproduce, including all the surface texture qualities, Sheen, Texture, and so forth? It's clear that the Realist has no clear, independent metaphysical reason for deciding how many attributes there are. We can't see them. What we see are the specific, determinate qualities. The abstract attributes are posited merely for the sake of explaining our possession of abstract concepts. The Realist is saying, in effect: our concepts are abstract; in order to be objective, they must correspond to something in reality that is intrinsically abstract; so for each concept there must be a separate abstract quality.

Now, this is the essence of what Ayn Rand called "intrinsicism," the idea that to be valid our concepts must reflect in a passive way a structure of properties that is out there and that can be grasped by a perceptual type mode of awareness. Realism is the quintessential form of intrinsicism. Now, I'm not going to go into everything that's wrong with that, but this is one of the fundamental points that Objectivism insists on, is that intrinsicism is wrong. Concepts can be objective; knowledge can be objective, without being passive, without passively mirroring what's out there. I'll be happy to expand on the reasons for that in the question period, if anyone wants to pursue it.

### Nominalism

All right, let me turn now to Nominalism, the other classical alternative in philosophy. Nominalists deny that abstractions exist outside the mind. They deny that any realm of entities exists beyond the world of particulars; in other words, they deny extreme Realism. And they deny that a particular thing's nature is bifurcated into a concrete and an abstract constituent. For the Nominalists, a thing is determinate through and through. The 18<sup>th</sup> century British Empiricist philosophy David Hume, who was a Nominalist, put it this way: the precise degree of a line, that is, the precise length of a line is not different nor distinguishable from the line itself, nor the precise degree of any quality from the quality itself. That is, a quality just is the particular, concrete degree of a quality, all right? If you ignore its specific character, there would be nothing left. So, the Nominalist does not analyze similarity into an element of identity and an element of difference. Two objects may happen to have the exactly the same determinate quality, exact same shade of red, but that would be rare. In the more common case, where the two objects are similar but not identical, they are just similar, period. Similarity is a metaphysical primary. It cannot be analyzed or broken down any further.

Well, if there's nothing universal or abstract in the world outside the mind, then concept-formation, the

epistemological process, cannot be a matter of direct awareness. The universality of concepts and the abstractness of concepts must be produced by some cognitive process. As John Locke, one of Hume's predecessors put it, features of general or abstract and universal belong not to the real existence of things – they're not out there in the world – but they are the inventions and creatures of the understanding, made by it for its own use. Well, how does the mind do this? By what process does a Nominalist think we generate abstractness and universality?

Well, as with Realism, Nominalism divides into two camps, one more extreme than the other. Extreme Nominalists regard concept-formation as a completely arbitrary process. We divide a thing's nature into attributes and group things together into categories any way we wish, and any way is just as good as any other. There are no objective constraints. I could group together a tree, a truck, and the city of Toronto and call them all "grucks," and I would have a concept. Now, of course, the only thing that those three items have in common is that I chose to call them by the same name, but that says the Nominalist is true of all of our concepts. The only thing that red things have in common is that we chose to call them by the same name, "red." Now, the obvious problem with this view is that it makes it impossible to identify new instances of a concept. Suppose I held up a new item and said, "Is that a gruck? I just taught you that, this new concept. Now, is this an instance of the concept, or not?" Obviously, you can't tell. Since I put together the concept by whim, you have no idea what my whim is going to deliver in the way of a judgment about this object. Nor do I. It's a completely arbitrary process. And the same is true, would be true, for all concepts on the Nominalist view. Extreme Nominalism is a clear case of what Objectivism calls the Primacy of Consciousness view. It regards concepts as wholly subjective, all right? And since all of our thinking depends on our concepts, the implication would be that thought as such is subjective. We're going to deal with a particularly striking case of this tomorrow in regard to induction. But let me leave the extreme subjectivist in his subjectivism for the moment, in his little world, and turn to the more reasonable version of Nominalism which is called "Resemblance Theory."

Resemblance Theory says that we group objects together on the basis of similarity, and we use a common name for the group. We identify new instances by determining whether they are similar to the instances that we now already, so the moderate Nominalist or the Resemblance theorist is trying to provide some kind of objective basis for our concepts. The classic example of this view is the one put forward by the third in the classical Empiricist trio, Bishop Berkeley, who came in time between John Locke and

David Hume. Berkeley said that, using the concept of "triangle" as his example, when we employ this concept and reason about triangles in geometry, we are not thinking about some abstract essence, Triangularity. What we're thinking about is a particular triangle, such as the one that we draw on the blackboard to serve as our demonstration. It's just that we realize that the conclusions we reach about that particular object apply to a range of similar objects, and we extend the same conclusion to other particulars, but without our ever having in our mind or seeing in reality some abstract essence, Triangularity.

And another version of this view has recently been developed in psychology by psychologists who have studied a phenomenon they call "typicality" or "prototypicality." What they've found is that, if you take a normal concept that people use for biological type or a type of artifact, "table," "chair," etc., some instances are more typical or representative of the concept than are others. Now, this may not seem like a very surprising fact to have spent all that research money on, but there are some interesting results from it, such as that, if you show people objects and ask them to classify them, they can classify the typical ones faster than the atypical ones. They can classify a beagle as a dog faster than they can recognize a Pekingese as a dog. And there are various other effects that this seems to have on aspects of cognitive processing. So, it's a real, as far as I can tell, it's a very real and somewhat interesting phenomenon. The theory that it has generated, however, is very much like Bishop Berkeley's. The idea is that a concept is organized around a prototype or prototypical instance. Plus, a Rule of Similarity that says, "Include anything similar to that prototype." So, the concept is "this thing, this prototype, this concrete instance, and everything that's similar to it."

Now, the problem with this theory emerges as soon as we look more closely at the role of similarity. First of all, it's not enough that an object by similar to the prototype in order to be included in the concept. It has to be sufficiently similar. If you showed me a demonstration in geometry regarding a particular triangle drawn on the board, and said, "And this is true of all similar objects," well, a square is somewhat similar to a triangle, but it's not similar enough, though, to fit the theorems. So, we need some way of delimiting the degree of similarity. An object has to have a certain degree of similarity to the prototype to count as an instance of a concept. This is not an insuperable problem. Various theories have been proposed for how we could mark off the degree of similarity, but it is a problem.

A more important problem is that when two things are similar, they are similar in some particular respect. If we don't specify the respect in which objects have to

resemble the prototype, then almost anything might be included in the concept. Suppose I try to teach you the concept of “red” by showing you a fire engine – at least, where I live, the fire engines are red – and I said, “OK, that’s a prototypical red thing. Now, include everything that’s similar.” All right, you go out and you see a garbage truck and identify it as a red object, because it is certainly quite similar to the prototype that I showed you. I haven’t given you the right information. I haven’t conveyed to you that you are to include everything that is similar to this object in the respect of color. But the problem is, how do we specify that respect? How do we represent the particular respect in which things have to be similar in the concept? I conveyed it to you by using the concept of “color.” So, we use concepts to identify a particular respect in which things might be similar, but when we’re forming concepts, obviously, we don’t have concepts like “color” available to us, so a theory has to be able to explain how someone who doesn’t yet possess concepts can individuate and isolate a single dimension of similarity. So a theory of concepts cannot merely invoke similarity to explain how we group things together. It has to explain how the mind isolates a particular dimension of similarity and a particular degree of similarity along that dimension. Because the Resemblance Theory treats similarity as a primary, something not subject to further analysis, it does not provide those explanations; and so the Resemblance Theory, despite its good intentions, tends to collapse into the more extreme, subjectivist type of Nominalism.

Before I leave the subject of Nominalism, I want to comment on one final point, one final issue, which I call the “analog-digital” issue. Now, for those of you who have worked with computers, I know these will be familiar terms, but let me explain them anyway. Think of the difference between analog and digital measuring instruments, such as a bathroom scale. On the analog scale, weight is measured by a needle that moves across a set of numbers. On a digital scale, the weight is represented by a number that appears and you read off the screen. On the digital scale, any weight that falls between, say, 164.5 and 165.5 will be represented by the same number, 165, assuming this measures only to the accuracy of a pound. On the analog scale, however, if you are exactly 165 pounds, the needle will be right on that mark. If you’re 164-1/2, it will be slightly to one side, and the same on the other side. So, an analog representation varies, it represents its different instances in a varied way; it varies along with the instances in correspondence to the things that it represents; whereas a digital representation uses a single item, in this case a number, to stand for all the differences, it ignores all the differences, among the things that it stands for.

Now, concepts, I believe, have both an analog and a digital aspect. On the analog side, we do make use of

typical instances; we do distinguish between the things that fall under a concept, between the typical and the borderline cases; and this has various consequences in our reasoning. But concepts also have a digital aspect. A Pekingese may be an atypical dog, but it’s still a dog. The statement, “This is a dog,” is completely true, it’s 100% true, same as it is for a beagle. As we saw earlier, a crucial part of concepts is that we treat the instances as identical. So, a theory of concepts has to explain both features. And let me give you a little example of both features at work in a genus-species hierarchy. If I showed you various pictures of robins, and I asked you which is a typical robin, you would have no trouble picking it out, and I think we’d all agree, the albino one is not typical, the one that’s two feet high is not typical. The typical one is little Robin Redbreast. So in that respect, we are alert to differences among robins. But now suppose I ask you, “which is a more typical bird, a robin or a penguin?” Well, again, you’d have no trouble answering; you’d say it’s a robin. But notice now you are treating the category of robins as a single unit. You are no longer taking account of differences among robins; you’re treating them as all identical, interchangeable instances of the term, of that category.

A more interesting aspect of the role of these two aspects of concepts, and it somewhat shows a little more why I’m drawing your attention to it, is in ethics. Consider the virtue of independence. Now, I could define this virtue in the abstract. I could state principles about the type of actions that it entails. And all of this is very important. But think how valuable it is to have, in addition, a clear example, an ideal case of independence, a prototype, a Howard Roark. There is a kind of thinking we can do, an application of the principle to our own behavior, that it helps us to have this sort of prototype for. That, indeed, is one of the functions of art, is to embody abstractions in concrete forms. Of course, Howard Roark is not an average instance, he’s an ideal instance, but nevertheless it’s still a prototype.

So that’s a way in which the analog aspects of concepts is very important in ethics. On the other hand, it’s not very hard to find examples of the danger of thinking exclusively in terms of typical cases or stereotypes. Racial and ethnic prejudice, for example, or tribalism, what Ayn Rand calls “tribalism” in general of this; it’s a phenomenon whose basis is at least partly epistemological in a kind of...well, she calls it “the anti-conceptual mentality. It reflects the way of thinking in which we see things as typical; we focus on what’s typical and we don’t grasp what is common; we don’t fully see the underlying identical principle that all human beings share. For example, we don’t understand...a typical tribalist doesn’t understand that people have rights qua individuals, qua human beings; and that’s regardless of racial or tribal or any other



specific characteristics, you individuals...all individuals have the same rights. [end of side A]

[beginning of side B] So, in that respect, it is crucial to think in terms...to ignore all the differences among people, to focus just on what is in common.

Now, as I said, a theory of concepts has to be able to account for both aspects of concepts, and that I think is one of the great virtues of Ayn Rand's theory, to which let us now turn.

### Rand's Theory of Concepts

Ayn Rand would agree with Realists that our concepts have a basis in reality. She would agree with Nominalists that what exists outside the mind is fully concrete, specific, and determinate. And she would agree with Nominalism that concepts are formed by noticing patterns of similarity among objects. But unlike Nominalists, she does not treat similarity as a primary, as a cognitive starting point. The starting point is difference. In perception, where concept-formation begins, we discriminate objects from their backgrounds. Discrimination means the awareness of difference. In vision, when we see a figure, we differentiate it from its background on the basis of differences in color, depth, and other attributes. So, if we were to observe two chairs, for example, our perceptual attention would be drawn to their differences. Each would be part of the background from which the other one is discriminated. The awareness of similarity arises when we perceive two chairs alongside something else such as a table. Each of the chairs differs from the table, the salient difference being in shape. Each of the chairs also differs from the other chair in shape, but the shape difference between the two chairs is not as large as the difference between either one and the table. It is this less than complete difference between the two chairs, which is brought out by the greater difference between either one of them and the table, that provides the foundation for seeing the two chairs as similar. That's the context in which we grasp the similarity between the two chairs.

Now, notice, without going any farther, that this solves the two problems of similarity that I mentioned before. First, there's no problem about determining how similar the two chairs have to be. They have to be more similar than either one is to the table. It depends on the context, but the context guarantees that there will be a certain limit within which...it guarantees...it gives us a standard for telling what is a sufficient degree of similarity. And secondly, there's no problem of isolating the respect in which the two chairs are similar. The respect in which they are similar, the respect in which they are perceived as similar, is determined by the respect in which they are differentiated from the tables. It's because we distinguish the chair from the

tables in shape, it is similarity of shape that our attention is drawn to between the two chairs.

Now, when we see the two chairs as similar, it's because the difference between them is less than the difference between either one and the table, the two chairs are seen as differing quantitatively from each other, whereas either one, either chair differs qualitatively from the table. If you look at the two chairs, you could see how one could be transformed into the other by merely quantitative changes, lengthening the back a little bit, lowering the seat, or whatever, whereas neither one could be transformed into the table by merely quantitative changes, it would take a qualitative change. [What about a wooden chair, a metal chair, and a metal table? Wouldn't it take a qualitative change in the wooden chair to morph it into a metal chair? And is morphing a metal chair into a metal table really a qualitative change any more so than morphing a large metal chair into a small metal chair? This example has its shortcomings or limitations.]

Similarity is, therefore, a quantitative relationship. It's a difference in degree, in specific measurement along some dimension. At the perceptual level, the dimension could be shape or color, tone in hearing, or warmth in the sense of touch, or whatever. But at higher levels of abstraction, this same principle applies. In forming the concept of "government," for example, we notice that certain social institutions differ qualitatively from certain other ones, such as the family, trade union, and church, or whatever, in that they use force. Force is the means by which they cohere, force as opposed to custom, tradition, contract, persuasion, and other types of interaction. But within the category of "governments," they differ merely quantitatively in the degree to which they use force, the scope to which their use of force intrudes into individual lives, the number of people who have control over the disposition of that force, etc. That is, they vary along quantitative dimensions.

The awareness of objects as similar, then, is the awareness of them as *commensurable*. Each could serve as a unit by which the others are measured, in the way a specific length such as an inch or a foot is a unit for measuring all lengths. All right, the measurements are relationships among the objects, quantitative relationships, each of which possesses and is seen as possessing its own specific degree on that dimension. And that is why she uses the term "units" to describe members of a group of objects on the basis of their similarity. It's as if we perceived this group of similar objects, or rather the relationships among those objects, in the form of a simple equation,  $b = ka$ , where  $b$  and  $a$  are two units, two members of the group of similar objects, and  $k$  indicates the quantitative relationship. For example, if this is a length of one foot, and this is a length of one inch,  $k$  is 12.

Now, in the perceptual level, we're not aware of the actual cardinal numbers, and indeed there are many dimensions along which we don't yet know how to apply a cardinal number system. We have to deal with ordinal numbers, first, second, third; we're ranking them just in a certain order of priority; but that is still a quantitative set of relationships. And so this little example illustrates the phenomenon. That is the way of symbolizing the form in which we grasp the relationships between similar objects.

All right, now, so far, what we have is a group of similar objects. We don't yet have a concept. The awareness of the units does not yet involve anything universal. We're dealing with a finite set of particular objects that happen to be before us perceptually, or that we can recall from memory. And we don't yet have anything that is abstract. We're aware of each object as having a specific characteristic, and we're aware of the specific set of characteristics, uh, specific set of relationships among objects in virtue of those characteristics. But those relationships that I've symbolized that way are the stepping stone to a concept.

### Concept Formation

We form a concept when we realize that any number of other objects could have any number of other specific measurements in relationship to these particular objects. We realize that the specific chairs can be ordered along a certain dimension, shape dimension, and that other items could occupy other positions on that dimension. In effect, we notice that the specific quantitative relationships that we've grasped in this form are simply special cases of a more general set of relationships that we can represent this way. Now, my apologies to the mathematicians for the bizarre look of this equation, but there's a method in my madness here. When we form a concept, we realize that this particular quantitative relationship, this between two of the units that we happen to perceive, is only an instance of a whole pattern of possible quantitative relationships that objects could have to a given unit. That is, if this is, again, length, we realize that an object could have any length and still be similar to the things that we have before us. And because an object could bear any quantitative relationship to a given object, any number of objects could be so related. So, we form a concept by expanding the set of concrete relationships that we grasp, unifying it into the idea of a single dimension on which any number of objects could have any number of places or specific measurements.

Again, I don't intend this to be taken literally that the child is thinking in terms of arithmetic and algebra. These are simply ways of representing the essence of the process. In Ayn Rand's words, "If a child considers a match, a pencil, and a stick, he observes that length is the attribute they have in common, but their specific

lengths differ. The difference is one of measurement. In order to form the concept 'length,' the child's mind retains the attribute and omits its specific measurements. Or, more precisely, if the process were identified in words, it would consist of the following: 'Length must exist in some quantity, but may exist in any quantity. I shall designate as "length" that attribute of any existent possessing it which can be quantitatively related to a unit of length without specifying the quantity.'"

Now, the specific example of length in her view sets the pattern for all concept-formation. We realize that membership in the similarity class, the class of units, can be expanded indefinitely by including anything that can be quantitatively related to units already included, on the principle that all members of the class must possess some particular membership, but may possess any. We recognize the commensurability of units without specifying the particular measurements, and the result of that process is a concept, [the result of] that process of measurement-omission is a concept.

A concept is a mental integration of two or more units into a new mental structure that allows us to treat those units as identical. Now, the specific units are still similar. They're not literally identical, as the Realist holds. But because their differences are differences in measurement, omitting their measurements is a way of ignoring their differences. Now, we don't ignore them entirely, and we certainly don't treat them as if they didn't exist. On the contrary, the specific measurements do exist, and we have to be aware of them in order for us to grasp the existence of a dimension along which things can be ranked. The point is simply that we grasp the dimension as something on which something can have any number of places. Indeed, since we are aware of the specific measurements, and because of that it's natural for us to choose a central or prototypical instance to serve as a representative of all the members. So, in this respect, her theory I think does account for what I've called the analog aspect of concepts, but it does also account for the digital aspect: we omit measurements on the principle that they have to exist in some degree, but may exist in any, and that is true equally and identically of all the units.

Now, once a concept is formed, it serves as a new mental unit represented by a word. The concept functions like a mental file folder, in which we retain information about things of a certain type. That's what allows the cognitive economy that I mentioned earlier. We reduce an unlimited number of numerically and qualitatively distinct objects to a single unit, and we can go on to group these new mental units together into higher-order concepts. For example, once we have formed the concept "chair," by omitting the specific measurements of chairs, we can go on to form the concept "furniture." We notice that the differences among tables, chairs, beds, etc., which we first

perceived as qualitative differences, are at this higher level of abstraction, can be regarded as quantitative differences. As against the qualitative difference all of them have to non-furniture items, like silverware, wall fixtures, etc. So, what's first perceived as a qualitative difference between tables and chairs is now seen as a mere quantitative difference in relationship to the qualitative difference between furniture and other artifacts. And in that way, as we expand the scope of our concepts and we reach increasingly high levels of abstraction, we are reducing quantitative to qualitative differences.

Now, in her book *Introduction to Objectivist Epistemology*, Ayn Rand goes on to derive a number of further consequences from her view of concept-formation, implications about definitions, about the contextual theory of knowledge, about axiomatic concepts, and many other issues. And, to my mind, as a philosopher, the appeal of her theory lies partly in its fruitfulness, its fertility in suggesting all these further implications. But I focused only on the foundations of the theory, because that's the first and the crucial test that a theory of abstraction, of concepts, has to pass. It has to deal with the basic problems that arise from the fact that concepts are universal and abstract. So, by way of summary and conclusion, let's compare her view with the two classical alternatives that I mentioned at the outset.

A concept, on her view, represents an abstract attribute such as length or a general type of thing such as chairs. The attribute or the type does not exist as such, as abstract, independently of a cognitive process of measurement-omission. What exists outside the mind are the particular objects with their specific relationships, qualitative and quantitative, and the potential that they have to be ranked in order along various dimensions. So, Ayn Rand's approach does not treat abstractions as intrinsic features of objects as the Realist does, but she doesn't treat abstractions as a subjective phenomenon either, a result of the mind grouping things together, more or less arbitrarily, on the basis of an undifferentiated awareness of similarity, as in Nominalism.

Abstractness is objective. It's a product of two facts, one about the world, and one about the human mind. It's a product of the metaphysical fact that the particular natures of things are quantitatively related along various dimensions, and it's a product of the epistemological fact that the human mind has the capacity to identify those dimensions by omitting specific measurements. And that's why, as I said at the outset, I regard this as a truly major contribution to philosophy, a solution to the Problem of Universals. Thank you. [Applause; end of lecture]

### Question Period on Concepts

OK, you all appear still to be awake, and we have some time left for questions, so I expect you have some.

[Male: *You said that similarity is not a primary, that difference is primary. Can that be analyzed any further?*] No. In the sense in which you intend it, I think the answer is no. I would regard difference as a metaphysical primary, that is, the existence in reality of a relationship of difference, that's a primary. Now, particular differences might be broken down into other differences. You know, the difference between capitalism and collectivism can be broken down into a number of component dimensions, but the phenomenon of difference as such has to be a metaphysical primary. And the reason I think that is that it's clear to me that the Law of Non-Contradiction has to be axiomatic, that it can't be a derivative conclusion, because it is presupposed by any process of reasoning that might generate a conclusion. The Law of Non-Contradiction has to be implicit in the Law of Identity. It has to be grounded somehow in the identities of things, but what Non-Contradiction says is that a thing cannot be A and not-A at the same, in the same respect. So, the implication to me is that the phenomenon of difference is implicit in the very existence of identities in the world, that part of what it is for things to have identities, which is the bedrock, inescapable fact, part of what it is for things to have identities is for there to be different identities in the world. Part of what it means, the essence of what it means to be something is to be not other things. So, difference is implicit there, and so I would regard it there as, well, in effect, as an axiomatic concept. Now, I want to emphasize that I am not saying that we can deduce similarity from difference. The relationship is one of the order in which they have to be grasped. Similarity has a positive nature. It's not simply the lack of a larger difference. You know, I said we grasp similarity as the smaller difference against the larger one. But similarity doesn't mean just the lack of a larger difference; it's a positive relationship of commensurability. It's an actual relation, numerical, quantitative relationship between two specific characteristics. It's just that, in order to grasp that, the cognitive context has to be prepared by first noticing certain differences.

[Male: *In pre-conceptual cognition at the perceptual level, how is it that essential defining characteristics of existents are grasped in order to form a concept, or the beginnings of a concept? How is it that the essence of something, that the essential defining characteristics of something are grasped without the concepts themselves?*] Well, typically, they aren't. That is, we normally, when we are...start forming concepts, we do so by grasping the more superficial similarities among things that are going to be instances

of a concept; and then, as we have the concept and learn more and more, we learn more and more fundamental aspects of those objects, and it's those that we then use to define...that we treat as essential and use to define the concept. For example, to use Ayn Rand's example, a child forming the concept of "man," of "human beings," "persons," would not begin by noting rationality versus other modes of cognitive functioning; it would begin by noticing shape and a certain sounds they make and a certain way of interacting with the child and so forth. It's only at a higher level of knowledge that the child can grasp a feature like reason. Now, does that answer your question, because... [Well, it sort of does, but it leads into...it's not related in terms of if the child is forming the concept of "man" by the perceptual evidence, and yet, what it is to *be* man, it isn't really a perceptual thing, I mean, rationality, as you say, if that's man's defining characteristic, and that's not perceptually obvious, yet that's really the key to forming or holding the concept of "man," so that when the child is recognizing that man is this sort of thing as opposed to the dog or whatever, or the table and so on. There's something about each being or existent that make them particular. It's not "qua man" that he's understanding them to be related and this sort of thing. So, I'm trying to understand how is he forming the concept of "man," if it relies on the non-perceptual similarity.] OK, well, be careful. This is a crucial point about concepts. Be careful not to identify a concept with the essential property that we use in defining the concept. The essential property is only one of all the properties common to the members of a concept, which we select for the cognitive purpose of giving a definition, of reducing a mass of information yet further, to a single proposition, OK? But that's not to say that the other attributes are unreal or are not part of the concept. They are a part of the concept, OK? So, what it means to be human is not just merely to be rational, but – or rationality plus animality – it's to have a language, it's to use technology, to survive by producing rather than by finding or taking what you need, to need art, OK, to be capable of laughter. All of these features that human beings have in common are part of the concept of "man." It's just that because reason is what underlies all of them, it's the best one to use as a definition, so that we can summarize all that information and condense it in the statement, "Man is a rational animal." But that does not exhaust the content of the concept. The content, really, is all the objects out there, with all of their properties. OK, OK. So that many concepts that...where the essential feature is one that can't be perceived. Rationality is one. Names of diseases, many of those are imperceptible, at least, essential properties; the symptoms are maybe perceptible, but the cause isn't. And when we get to

social and economic systems, you know, we're way beyond the range of what can be grasped perceptually.

[Male: *Dr. Kelley, Ayn Rand said that a concept is...that in developing a concept, we don't...we isolate certain attributes as fixed across a particular concept and allow other attributes to vary...squares, length was allowed to vary. Are there not also other concepts where length is the fixed attribute and it's other things like shape that vary...say, in a clothing store, for example, everything that is on the rack that is called "small," the length is the fixed attribute, and it's...colors, shapes, and other attributes vary...*] Right, OK, that's an interesting point. When we classify objects, we normally try to classify them by essential principles. So, if I said "classify clothing," since clothing is a human artifact, the essential property is going to be its function, so you would classify it by the different functions they perform, such as covering what part of the body, and then what design intended for what circumstance, so you classify into coats, jackets, trousers, blouses, skirts, etc. But of course, you also can classify by other attributes, such as color, length, style, and a clothing store, to serve the function of serving its customers, would very naturally classify by what in a generalized context would be non-essential, but in this context is essential, namely, length, because a given shopper doesn't care about the other lengths. But notice that classification is not the same thing as concept-formation. Concept-formation begins by classifying things together into a group, but it then has the other step of measurement-omission that produces a new concept, that is, a new mental unit, OK? Now, once we have a stock of concepts, we can group things together conceptually and not just by perceptual similarity. We use concepts to classify things into various categories, but unless the category is of fairly significant importance to us, we don't form a new concept for it. That is, we don't have a concept for clothes that are 42 Normal, OK? We don't need it; a description will suffice. This is what is sometimes called "Rand's Razor." Don't multiply concepts beyond necessity. There are many classification tasks that you can perform to serve the needs of a moment or a specialized context where you don't need a new concept, and therefore you don't form one. Forming a new concept is a major epistemological event. It's like having kids. [Laughter.] It's not to be done lightly. OK, so, it's perfectly rational to do...to proceed as you say, but I would say, not to form a concept. [Same audience member: But there may be a case where length is the unifying aspect.] Oh, yes, yeah, certainly.

Allen? [Male: *Yeah, it follows that question. Is there a criterion that can be easily stated for when a concept has to be formed, as against a description? I know one condition would be that you actually condense*

a description into a mental unit that's useable. So, it occurs to me that, suppose the trouser shelf...and the people in the shop and in the business have reduced to a description, you know, whatever, 30 extra wide has the "smalls," and the small trousers are over there. Now, there's a mental unit which certainly can be condensed for economy purposes, but it's not clear to me whether that becomes a new concept or not... Well, it seems to me that it wouldn't because, even though you have reduced a description to the single word "small," you've been able to do that only because the context makes it clear that what you're talking about are small trousers. If I came in, and I said, "This is small," all right, this is small, even though I can't put it on, all right, so this is not an instance of the single word, term there. In other words, what you're doing is taking a word that has a concept you've already formed and simply...this is clearly, I think, an implicit description, and the evidence for that is that the term "small" is still used, still available for...outside that specialized context. I think if...to...you'd really need a new word, vocally. [Same audience member: That's why I said "smalls," because I'm...that if there is a concept... subdivision of the concept "trousers," and the question I pose is: with regard to subdivision of concepts, where you've got a structure generated by the concept itself, you're bound...I mean, I'm sure that many subdivision concepts began as...would be conclusive.] All right, I mean, if we actually got into some of the etymological questions, yeah, surely there are cases where...we all should be able to think of lots of cases where a word that was first a generalized name for an attribute, like an adjective, and it's nominalized and becomes a noun designating a specific type of thing. I don't know why I can't think of an example. [Male audience member: *blondes*.] All right, sure. That seems to be sort of like a borderline case. [Laughter.] It still seems to me pretty clearly short for a blonde person or, more normally, a blonde woman. But, yeah, there's too many other cases like that, and I'm not sure that there is any hard and fast criterion. You just have to look at how...at the whole context of a person's knowledge to see whether...how fully this has crystallized into a new unit. And I think there will probably be some perceptual tasks that would allow you to do some experimental work on it. For example, there are experimental findings that suggest that people have access to a concept faster if you prime them by using other concepts in the same area, the same genus, for example, or giving them the genus name, OK? So, I'm going to have the word "beagle" come faster to my lips, if – you know, ready to identify objects – if you've already primed me for the category "dog," OK? So, you might use that kind of experiment to decide whether – you know, give a person the word "trousers" and then see whether "small" comes faster in

a way that is different from "small" in other applications, all right?

[Male: *Wouldn't that be called "abstractions from abstractions"?*] That would be...the phenomenon we were just talking about? [Audience member: Yes.] Yes, it would be, right. "Abstractions from abstractions" refer to all those concepts – which includes, in fact, the vast majority of them – which are formed on the basis of grasping similarity in a conceptual rather than a purely perceptual form, whether they are wider integrations, narrower differentiations, extended special case concepts, or whatever, so, yeah, that would be an abstraction from an abstraction, OK?

[Male: *What is more fundamental, induction or concept-formation, a process of induction or concept-formation?*] Well, if by "induction," we mean the specific process of supporting a general proposition, "All S are P" or "No S are P," the process I'm going to talk about tomorrow, I would say concepts are...concept-formation is more fundamental, because you need the concept to indicate the category of things you're going to talk about and the attribute you are generalizing that they all have. Now, there are processes that we could call broadly "inductive," that work along – in fact, concept-formation itself could be described, as Ayn Rand herself describes it, as in essence an inductive process. That's why induction and concept-formation are so tightly tied together. They're both processes of starting from particulars and extracting from them a general rule or principle. And so, in that more generalized sense, induction is something that happens along with perception {does he mean concept-formation?}, and I would even say, in that broader sense, if you try to explain animal functioning, all right, there are various processes of association that they have that might, in a sense, be called "induction," at least for animals that are capable of more sophisticated types of learning, such as the primates and so forth. But in the narrower, logical sense of an inference to support a proposition like "All S are P," you need concepts first.

Eric? [*Somewhat inaudible question about birds, robins vs. penguins to native of Antarctica, the latter are typical.*] I'm not sure about the determinants of typicality, and I'm not sure that's true, what you've just suggested is true. I would be surprised if familiarity had no effect, but I would also be surprised if it was the only factor, because typicality is – in many cases, anyway, at the perceptual-level concepts – the typical instance is one that has an average value on the relevant dimensions, or roughly an average value. That is, it's midway between the borderlines on the relevant dimensions. It's got kind of an average shape, average size, average color along the color-shape-and-size

dimensions for a given kind of bird or whatever it might be. And a penguin just is... there's nothing you can do to put it in the average category. I mean, it can't fly. It's got zero...it spends zero amount of its time flying, and that's not going to be an average. More likely, for people who spent a lot of time and whose livelihood was deeply involved with penguins would probably have a separate category for penguins and might not even...they would just...would not think of them as birds in the first place. There would be penguins and these flying objects – there were penguins and these other species, and somewhat like the way children think about human beings and animals. Kids don't think about man as an animal. Human beings are like us, and animals are everything else, starting with dogs and kind of shading off down into lizards and locusts and so forth. But that's a guess. I don't know all the data on that.

[Male: *I still don't fully understand your answer to...Where's the difference between someone saying there's a difference between a robin and a penguin, but what about a robin and a blue jay?*] Well, they're both typical. They both have...people actually do studies where they've come up with a number system for ranking typicality. [Laughter.] [Audience member: *That's what I don't understand.*] *Blue jays and robins are both real high...*[Laughter.]...*I mean, for all I know, blue jays may be higher.* So, yeah, but you know, that's actually interesting, the lists from blue jays and robins down to wrens and tufted titmice and, I don't know, owls maybe, and then down to, you know, the really weird things like auks and dodos and ostriches and so forth, penguins. What these are done...these are compiled from judgments that subjects make, and what you do is to take all the answers and compile them, it's like a certain voting system where you decide who wins by compiling all the first, second, third-order preferences.

[Male: *(Somewhat inaudible question about the voting referred to above).*] I think that's true. I think that would be a factor, yes, I certainly agree, and I think there's evidence that familiarity and frequency of experience with something has an effect on typicality. I'm just not sure it has the only effect.

[Male: *(Somewhat inaudible question about resemblance theory).*] I'm going to have to reformulate to see...just make sure I understand it. I claim that the... Ayn Rand's theory avoids the problem that the prototype theory has, namely, of...she explains...the prototype theory has a problem of saying when two objects are sufficiently similar to be counted...when an object is sufficiently similar to the prototype to be counted in the concept, and she solves that problem, because you're always dealing with a contrast object, like a table. Here's where I'm not sure I get you.

If...couldn't a prototype theorist say, "That just shifts the problem a little bit, because now you have to explain when two contexts are sufficiently similar to be counted as...two contrast objects and the contexts they determine are sufficiently similar for two people to be considered as forming the same concept. Is that a fair...OK. Well, in a way that's...I would not regard that as a fair question, because it's comparing two...questions at two different levels. One is: what process do we think the child is following? That's where the objection to the prototype theory is located. And the other one is: how do I, as an epistemologist, defend my own conceptual practice in saying that two children are forming the same concept? That's a meta-level question, which presupposes that we already have an enormously complicated context, which gives us resources to answer that, such as, we know that...we fully possess the concept "chair," so we know what the borderlines are; we know that concepts involve definitions. So, if you ask me how would I establish that two children are forming the same concept, I would look at a much wider range of issues. In other words, I would ask...I would test them on things: do they recognize this as a chair, even though it's not rectangular? Can they give me any kind of definition of a chair? Can they locate it in the right higher-order concept, that's "furniture"? And I would expect that there would be some differences between...particular contrast object that they might have used at first in order to form the concept. Indeed, I would expect that...learning concepts, conceptual learning, is largely a process of...where people start at somewhat different levels. People include in a concept somewhat different ranges of things, and then gradually through learning, the concepts as they...especially as they learn language and talk to people, concepts come closer and closer to having the same content. Children notoriously over-extend their concepts. The child looks at the moon and says, "That's a ball," OK? Well, he doesn't have the concept of "round," so what else can he call it? But eventually he'll learn that that's actually not a ball, that that's a moon.

All right, we have to end here... [Applause.]

## **Part II: Problem of Induction**

Dr. Kelley: [Beginning of lecture not on tape]...to describe various types of inferences and, even more broadly, various types of cognitive operations, but today I'm going to be concerned with a specific type known as inductive generalization. Inductive generalization means inferring that something is true of an entire class of things, because it has been found to be true of certain members of that class. We observe that certain members of the class S have a certain property P, and we draw the conclusion that all S's are P. Copper melts at 1083 degrees Centigrade; that's true of all copper. All men are mortal. In both cases, we support these propositions by the observation of some particular pieces of copper and some particular human beings, all the ones who have lived so far. In the language of statistics, inductive generalization is drawing an inference from a sample to a population as a whole, from some S's to all S's.

Now, at its root, induction relies on causal connections. What justifies us in assuming that all copper melts at 1083 degrees Centigrade is the existence of a causal link between the molecular structure of copper, which makes it an instance of copper in the first place, and the temperature at which it changes its state. Whenever we draw an inference that all S are P based on a sample, we are relying on the assumption that S and P are connected and that connection, in one way or another, is a causal connection.

Now, here's the problem. When I generalize, I start with information about certain members of a class, and I draw a conclusion about all members, most of which I have never observed and never will. How can I be sure that what's true of some members is true of all the others, the ones I haven't observed? Am I not going beyond my evidence? Isn't there a leap of faith here? If I say that the inference is supported by a causal connection, well, that only shifts the question. How do I know that there is a causal connection between S and P? Even if S and P have always been connected in the past, how can I be sure that that connection will persist in the future? The sun has risen every morning in recorded history, but how can I be sure it's going to rise tomorrow?

Now, you can see that this is a serious problem. If we can't rely on induction, then we can't employ deductive reasoning either. Deductive reasoning is the application of a general principle to an instance that falls under it. All men are mortal. Socrates is a man. Therefore, Socrates is mortal. But such an inference is only as good as its premises, and if the premise is a generalization, then, unless we can support that

inductively, the rest of the inference is useless. In fact, if we cannot generalize inductively from experience, then there's no point in forming concepts in the first place. The point of forming concepts, of classifying things, is to discover the attributes and ways of acting common to all members of a category, so that we can be prepared to deal with new instances as we encounter them, and we can't do that without inductive generalization.

My goal today is to give an answer to the Problem of Induction. Now, this answer is not original with me. Some people, when they heard that I was going to be giving these lectures, said, "Oh, are you going to solve the Problem of Induction? That would be nice." And I felt a little embarrassed because I thought it had already been solved. The problem is not even...the answer, anyway, is not even really original with Ayn Rand. I believe it's implicit in a broader Aristotelian tradition in philosophy, although we're going to see that her theory of concepts allows a much stronger formulation of certain points, a much more impregnable defense of the Aristotelian position than the actual Aristotelians could offer.

Now, we need to look at this problem at two different levels. First, we need to ask how to validate the Law of Causality, the fundamental principle that says reality is governed by causality in the first place. The Law of Causality, as we'll see, is an indispensable metaphysical basis of induction. Secondly, we're going to examine certain questions about how we establish specific causal connections, such as the ones I've used as examples, the melting point of copper and so forth. These are questions about the methodology, the epistemology, the logic of induction. And we'll see that, at this level, particularly, a good theory of concepts is indispensable.

### **Validating the Law of Causality**

All right, now, I want to begin by taking a closer look at the problem, and I want to do that by reviewing for you the arguments of the man who first posed this problem in its most powerful and dramatic way. I'm referring to David Hume, the 18th century British empiricist. Hume believe, these are his conclusions, that there are no inherent connections between causes and effects, between events that we see to be...that we observe to follow one another in reality. They do follow one another, but there is no inner connection between them, there's no necessity. In his own words, Hume put it this way: All events seem entirely loose and separate. One event follows another, but we can never observe any tie between them. They seem conjoined but never connected. Therefore, we

have no rational justification for generalizing beyond the range of our immediate experience. The fact that we've seen the sun rise every day up to this point gives us no rational basis for expecting it to rise tomorrow. The fact that bread has nourished us in the past gives us no reason for thinking that it will not be toxic the next time we have a slice. The regularities that we observe represent a kind of metaphysical coincidence, a coincidence like the fact that for the last 70 years Russian agricultural areas have had unusually bad weather. [Laughter.] But because there are no inherent connections among events, there is no reason to expect this coincidence to continue.

Now, the problem is not that some unpredictable factor might interfere. The problem with knowing that the sun will rise tomorrow is not that there may some comet that's going to collide with the earth and disintegrate it or at least stop its rotation. No. Even if there...without any interfering events, we can't know that the sun will rise tomorrow, because we can't know that it is in and of the nature of the earth in its present state to continue rotating. Even if no further unpredictable event occurs, we don't know that causality will continue in its normal course.

Now, of course, we do...we expect that it will. We expect things to continue acting as they have in the past. We can't help ourselves. But our basis is not reason. This expectation is a kind of psychological habit that has been induced by the regularities that we've observed. We're like superstitious people who have observed some chance coincidence, breaking a mirror and having a run of bad luck, and we expect it to continue. All of our beliefs about causality, says Hume, have the logical status of superstitions. They're real enough as psychological phenomena, but they have no objective basis.

Now, why would anyone believe this? Well, says Hume, imagine you had never seen water before. Could you tell, just by looking at the water, observing its fluid and transparent character, that you could not breathe under it, that it would suffocate you if you held your head under it? If you had never seen fire, could you tell, just from the observable color and shape of the flames, that it would burn you? No. Cause and effect are distinct, and observing one by itself will not allow you to deduce the other.

Very well, you might say, but once I have observed what effect follows from a given cause, I can see that they're connected, and then I can generalize that other instances of that cause will have the same effect. Once I observe that the flame burns, then I know that future flames will burn. But wait a minute, says Hume,

can you really see the connection? You see the flickering yellow flames of the fire, you feel the pain when you put your hand in it, but do you in addition see a connection, a relationship of necessity between cause and effect? No, again. A sensation of yellow is followed by a sensation of pain, but all there is is the temporal sequence in your experience. You observe a billiard ball colliding with another one, you observe the second billiard ball moving off, but, again, all you observe are the two events; you don't observe, over and above those two events, a relationship of necessity connecting them.

Now, to reinforce the point, Hume offers us two additional arguments. First of all, even after you have seen the second billiard ball roll off in a certain direction, couldn't you just as easily imagine it rolling off in a different direction, or jumping off the table, or collapsing into a puddle? In his own words, even after we have observed an effect, "the conjunction of it with the cause must appear arbitrary, since there are always many other effects which, to reason, must seem fully as consistent and natural." Secondly, if you did observe some real connection between cause and effect, then you would be able to generalize from a single instance. After all, in mathematics, in geometry, if you observe a relationship...if you follow a proof from a single illustration on the board that the internal angles of a triangle equal 180 degrees, then you can see by the internal necessity of that demonstration that the same would be true of all triangles. By parity of reasoning, if there's a necessity between cause and effect, then observing one instance of that relationship should allow you to form a generalization without the need for repeating the experiment in any way. And yet, Hume says, it is a characteristic of induction that we do have to repeat. We have to gather a series of examples of the same thing in order to support a generalization. Now, please remember this point, because I'm going to come back to it later on.

Now, at this point, you might try a different tack. You might appeal to the Law of Causality. And for our purposes, we can think of the Law of Causality as having two basic elements, two clauses which I have stated on the handout. The first clause says that everything that happens has a cause. Events do not occur randomly. The occurrence of an event and its particular nature are determined by the circumstances in which the event occurs, including the nature of the entities that act, the preceding events, and so forth. Secondly, the Law of Causality says that the same cause has the same effect. Once we have discovered the factor responsible for some event in a particular case, we can be sure that the same factor will have the same



consequence in any other case, so long as the other circumstances that are relevant have not changed.

So, you might answer Hume as follows: True enough, you might say, I don't actually perceive causal necessity. I see the fire, I feel the pain, but I don't literally perceive a relationship of causal necessity between them. The Law of Causality, however, tells me that such a relationship exists. The first part of the law says that the pain must have some cause, and once I've isolated the fire as the causal agent, the second part of the law tells me that the fire will have the same effect in any other instance.

But now, Hume will ask the obvious question: What justification do you have for believing in the Law of Causality? How would you validate the Law of Causality? You can't establish it by inductive reasoning, he would say, because that would be circular. Inductive reasoning presupposes the Law of Causality, so you can't turn around and use inductive reasoning to prove the Law of Causality; that would be circular. So, the Law of Causality would have to be a self-evident axiom, or at least a principle that could be deduced from some axiom.

But, says Hume, the mark of an axiom, of a self-evident truth, is that we cannot deny it without contradicting ourselves; and the test of whether a proposition is contradictory is whether we can imagine it. We cannot imagine a case in which A equals not-A. That's why the Law of Identity, A is A, is self-evident. It cannot be denied without contradiction; we cannot imagine it being false. But we can imagine that fire soothes on Monday, Wednesday, and Friday and burns on Tuesday, Thursday, and Saturday – Sunday being a day of metaphysical rest. [Laughter.] The Law of Causality can be denied without self-contradiction, so it cannot be an axiom. And even if it were an axiom, finally, that wouldn't do us any good. Axioms, says Hume, are merely relations among our own ideas; they are arbitrary constructs that we make true by the way we define our terms, and therefore they don't tell us anything about reality.

So, Hume's conclusion, to summarize, is that all events are "loose and separate," to use his phrase. They follow one another in a temporal sequence, but they are not connected. Any event can be followed by any other. We may notice certain patterns of regularity, but these don't reflect any inner connection between one event and another, or between an entity and the way it acts. So, when I turn on the stovetop burner on my stove, I expect gas to come out the jets. It's done that a thousand times before in the past, but it would be mere coincidence if that's what happened this time. It's

equally possible and, indeed, equally likely that some other, any number of other events could have happened. The stove could explode. It could rise up off the floor. An army of ants could come streaming out, marching in unison and singing "Deutschland Uber Alles." [Laughter.] Anything is equally possible.

Now, this is a bizarre view of reality. It's a kind of nightmare vision. My image for it is like spending your whole life driving on ice, which we've all had some experience with lately. [Laughter.] And frankly, it's incredible. But it's also been very influential in philosophy. So, let's turn now to the opposing case and see how to answer Hume's argument. How can we validate the Law of Causality in the face of his assault?

Well, Hume was correct, first of all, in saying that the Law of Causality cannot be established by induction – that is, we cannot proceed by proving first that water obeys causal law, then that fire obeys causal law, then that the economy obeys causal law, and then generalizing that all phenomena obey a causal law. Induction presupposes the law; we need the law in order to establish any specific conclusion, so the Law of Causality must be axiomatic.

Now, most philosophers have held that if a statement does not rest on induction, then it does not derive in any way from the evidence of the senses. In their view, induction is the only way to bring empirical evidence to bear on a general statement such as the Law of Causality. The implication is that an axiom must be a statement adopted by reason operating in some way independently of the senses; it must be a priori. And Hume, like many other though not all philosophers, drew the further implication that axioms must be subjective and must be arbitrary; we adopt them by convention rather than by discovering the facts. After all, if perception is our only contact with reality, as the Empiricists rightly believe, and if induction is the only way to extract information from perception, then a principle not based on induction cannot be true by virtue of some contact with reality. It must be true merely by convention; it must be arbitrary at its root.

Now, Objectivism, along with the entire Aristotelian tradition, I might add, does not accept this argument. An axiom is a self-evident statement, and a self-evident statement is one whose truth can be observed directly. Such a statement formulates in explicit terms what is implicit in perceptual awareness. Inductive generalization is not the only way to extract information from the senses. So, let's turn our attention to perception and see how we can try to extract the Law of Causality from what is given.

OK, the essence of perception is the awareness of entities, i.e., of objects as wholes, discriminated from their background, existing and acting as units. In this respect, perception is different from sensation, which is a more primitive level of awareness. A sensation is a momentary awareness of an isolated sensory quality, such as a patch of color, a feeling of pressure on the skin, a taste of salt. A perception, by contrast, is the perception of an entity as a whole, of a constellation of qualities which are integrated.

Now, of course, we are aware of qualities in perception, but we're not aware of them as isolated, atomistic units. Nor do we perceive the entity as a bare framework on which these qualities are hung, like ornaments on a Christmas tree. When we see an apple, we are aware of its color, its shape, its size, its texture, and all its other qualities as integrated elements of a single identity. We perceive the object as a unified nature, and the object is its nature. This is the basis for the Law of Identity, the axiom which says that an existent is something specific. Whatever it is that we perceive, it is something specific; it has an identity; it is its identity; it is what it is.

Now, we can also perceive the actions of entities, the way they move and change. And in that respect, too, perception differs from sensation. Sensation is a momentary experience. Change at that level is experience as the substitution of one quality for another, as in a kaleidoscope. But perception does not consist of isolated snapshots in that way; it's the continuous awareness of objects over time. A hunter follows the prey as it moves across the field. A doctor listens to his patient's heart as it beats. Both of them are continuously aware of a single entity as a single entity and aware of its actions, and they're aware of the actions *as* the actions of the entity. Whenever we perceive an action, we perceive it as the action performed by something, just as we always perceive attributes as attributes of something. There cannot be motion without a moving object; there can't be a dance without a dancer. But when we form a concept for a type of action, of course, the entities involved become omitted measurements. A waltz is a waltz, no matter who's doing it. But the measurements must exist in some form. Without something that acts, there cannot be any action.

Now, all of this, I claim, is implicit in perception, and the Law of Causality is simply a way of making it explicit, putting it into explicit, conceptual form. And we can do this in three steps, although I want to emphasize as strongly as I can that these are not steps of an argument. Since what I'm presenting to you

here is axiomatic, I would be undercutting my own case if I presented it to you as a proof. It's simply a way of drawing your attention to what is implicit in your own perception so that you can grasp for yourselves the connection. In psychology, I've heard reference to something called "guided fantasy." Well, this is sort of the opposite. This is a guided perception.

All right, step one. The first part of the law says that all actions have causes. To say that an action has a cause is to say that it is produced by something, that it is a dependent phenomenon, not a self-sufficient existent, not a primary, and that is precisely what perception tells us about actions. They do not occur without entities. It may be, as Hume says, that we can imagine an action as a self-sufficient existent occurring without a cause, but if so, we can do that only by dropping the context of what we observe perceptually. If you examine your perception, if you are anything like me, you cannot imagine an action occurring without an entity, and so long as you hold the perceptual context, you literally cannot imagine it. It's like...you can omit the measurements, of course, when you form a concept. But just as you can't imagine a color that is red but not any specific shade of red, it's impossible to imagine yourself observing, so long as you're in touch with your perceptions, an action occurring with no particular thing that acts. OK, so the first part of the law is simply a way of identifying the observable fact that an action is the action of some entity. It identifies the dependent nature of the action, the fact that actions are not primaries, metaphysically.

Step two: If actions depend on entities, then an action must depend on the nature of the entity that acts. A thing is its nature. If we try to imagine an action that depends on the entity but not on its nature, we have to imagine the entity as something distinct from its nature. We have to drive a kind of metaphysical wedge between a thing and what it is. But a thing is what it is. There's no such gap.

So, now we know that an action must depend on the nature of the entity, and we are ready for step three, the final step, which is to realize that, because an action depends on the nature of the entity that acts, an entity with the same nature, operating in the same circumstances, must act in the same way. The same cause must have the same effect, which is part two of the Law of Causality.

Now, it's true, as Hume said, that we can imagine the same cause giving rise to different effects, the same entity acting differently. But we can do this, once again, only by dropping context, in this case, what we have just acknowledged in step two, that the action

depends on the nature of the entity. Uniformity is implicit in the idea of an action depending on the attributes of the entity or of those attributes producing the action. If you take away uniformity, you take away any concept of “production” or of “dependence.” Suppose I said the motion of a billiard ball when it is struck is produced by its rigidity; it’s that rigidity that makes the ball move when the first ball hits it. But now, suppose I say that another ball with the same rigidity, or the same ball at a later time, might stand still when it is struck, or collapse on the table like a rotten mushroom, or recite the Gettysburg Address [Laughter.], and that in all those cases, it would still be the rigidity that is causing the action. Well, that’s incoherent. By imagining these different effects, I have severed any connection between the rigidity of the ball and the way it acts. I have, therefore, withdrawn the admission I made previously that the action depends on the nature of the ball. A cause that does not act uniformly is not a cause. We’ve replaced the connection between the entity and its actions by a random relationship, and a random relationship means there is no relationship.

Well, that’s it. That’s the validation of the Law of...that’s our guided perception. To summarize what we’ve seen, let me quote from a work by an Aristotelian logician, H. W. B. Joseph, who was an Oxford philosopher writing in the earlier part of this [the 20th] century. The passage I’m going to read to you is from *Introduction to Logic*, a chapter called “Presuppositions of Induction.” Joseph says:

“If a thing A under conditions C produces a change X in a subject S [it gets easier from here on][Laughter], the way in which it [A], the way in which it acts must be regarded as a partial expression of what it is. It could only act differently if it were different. To say that the same thing acting on the same thing under the same conditions might yet produce a different effect is to say that a thing need not be what it is, but this is in flat conflict with the Law of Identity. A thing, to be at all, must be something, and can only be what it is. To assert a causal connection between A and X implies that A acts as it does because A is what it is, because in fact it is A. So long, therefore, as it is A, it must act thus, and to assert that it may act otherwise on a subsequent occasion is to assert that what is A is something else than the A which it is declared to be.”

In short, the Law of Causality is the Law of Identity applied to action.

Events are not “loose and separate.” They are tightly constrained by the entities that act. And once we have identified specific causal connections, our confidence that things will continue to act accordingly

in the future is not groundless or subjective. It’s a recognition of an objective fact. So, we have answered Hume’s argument. We’ve rejected his assumption that axioms are arbitrary conventions. They’re grounded in perceptions. We’ve rejected his assumption that what we can imagine is a valid test of what is possible. We can imagine violations of the Law of Causality only by dropping the context of what is given in perception. Our imagination may be interesting from a psychological standpoint, but it has no epistemological significance.

There’s a third point that I didn’t emphasize, but I think it deserves some mention. Hume regarded causality as a relationship between events, whereas Aristotelians regard causality as...primarily as a relationship between an entity and its actions. This is a metaphysical difference in viewpoint, but it has an epistemological basis implicit in what we have just been over. The basis is that Hume did not see perception as the foundation of knowledge, the epistemological starting point. For him, the starting point was sensation. What is given, in his view, are merely isolated qualities in momentary acts of awareness. Perception, said Hume, is the result of putting those sensations together by some process of association; and the entity that we perceive is simply a set of qualities, but we don’t see any metaphysical glue holding those qualities together. The entity is a constructed invention of the mind. So, we have no reason for believing that all attributes are attributes of entities, nor that all actions are actions of entities. So, Hume deprived himself of a crucial point on which the Law of Causality depends and, therefore, could not see any reason why events could not be primary and autonomous types of existents. The dance is given, but not the dancer.

Now, of course, Hume recognizes that, in some cases, we do associate an effect with an attribute of an object, as well as with a prior event. We associate the motion of a billiard ball with its rigidity, as well as with the prior event of the collision. So, Hume acknowledges the facts that Aristotelians would cite in supporting their view of causality, but the form in which we’re aware of those facts, says Hume, is the form of sensations. So, the rigidity of the billiard ball is not, as it’s given in experience, an enduring property, a part of a nature; it’s a momentary content of a momentary act of awareness. Even an attribute, at the level of sensation, is experienced as an event; it’s a momentary content of awareness to be replaced a moment later by something else. It’s the occurrence of a certain quality at a time and place. So, you see, the metaphysical difference is not fundamental; it’s the epistemological difference starting with sensations vs. perceptions that is crucial here.

## The Method of Induction

OK, well, let me turn now to the second major topic to be covered today, which is the method of induction. I think it's clear that the Aristotelian and the Humean philosophies differ radically in their view...basic view about causality and the Law of Causality, and the same is true about their views of the procedure by which we establish particular causal connections. So, let's turn to the second broad topic, and let me start with a few remarks about Hume and the Humean tradition in inductive logic.

Hume and his followers, as we saw, do not see any interconnection between cause and effect. There are observable regularities, and antecedent factor A may always be followed by a certain type of effect E. They may be constantly conjoined, in Hume's phrase, but there's no reason why A must produce E. So, observing a causal regularity, on their view, is like drawing colored balls from an urn, to use a favorite example of statisticians. If we draw a series of balls, and all of them are black, we may begin to suspect that all the balls in the urn are black; but since there's no reason why they must be black, we are limited to a claim of probability. We can't be sure that the next one we're going to draw is black. The more balls we draw, the more probable it becomes that they're all black, but we can't be certain. The probability never reaches 100% until we've drawn all the balls out.

Now, in the same way, when we put our finger near a flame and feel heat, we observe an association between flame and heat. But because there is no inherent reason why flames must be hot, a single instance is not very good evidence for a generalization that all flames are hot. We have to repeat the experiment over and over, getting more and more confirming instances. The more instances we have, the higher the probability rises that all flames are hot, although in this case, since we cannot ever observe all flames, it being an open-ended concept, the probability never reaches 100%. So, in this way, Humeans typically regard inductive logic as a branch of probability theory, and they hold that the strength of a generalization depends in an essential way on the number of confirming instances.

The Aristotelian approach is quite different. The Law of Causality tells us that the same cause has the same effect, so once we have isolated the cause of a given effect in a given case, we can generalize immediately. We know that that cause will always have that effect. The problem is to isolate the relevant causal factor in a particular case. When we observe an action

or an event, we know that it depends causally on the natures of the entities that are present in that situation, but we don't know which aspects of their identities are relevant. We need to discover which particular attribute of a thing gives it the capacity to act as it does. The fact that salt dissolves in water tells us that salt is soluble, that solubility is part of its nature, but we don't know what particular attribute of salt makes it soluble, gives it that capacity for action. This is what we need a method of induction for. And that method, which was first proposed in an early version by Francis Bacon, and later refined and systematized by John Stuart Mill, is the familiar method of systematically varying the factors experimentally. So, let's look briefly at three of Mills' methods. These are called Mills' Methods of Induction. He named five, but these are the three primary ones, and I've illustrated them on your sheet, on the handout.

You'll notice that in all three cases, the Method of Agreement, Method of Difference, Method of Concomitant Variation, on the right-hand side is the letter E, which stands for the effect. Preceding E there ought to be an arrow. I didn't think...my printer doesn't do this kind of thing, and I didn't think to put it in before I started Xeroxing, so I left it to you to do. Please put in arrows...arrowheads on the right-hand side of all the dashes.

In the Method of Agreement, we look for a unique factor that is common to all the cases in which the effect occurs. We hold one factor A constant, and we vary all the others. For example, we want to explain why all the people who ate a restaurant on a given night got sick, and we identify, we find that the common factor is that they all had the Brussels sprouts. So, we have isolated factor A as the cause, because it is the only factor that satisfies the Law of Causality, in particular which says if a given factor is sufficient to produce a given effect, it must produce that effect even if other non-relevant factors are altered. If in case #3, under the Method of Agreement, we had not obtained the effect, then we would know that A was not the cause, because the same cause must have the same effect.

Let's look at the Method of Difference. In this method, we keep all the factors constant except for the one we are testing; and if the effect occurs in the presence of A, but does not occur in its absence, we know that A is at least a necessary condition for the effect and that none of the other factors are sufficient. The Law of Causality tells us that if B were the cause of the effect, if B were the cause of E, in case #1, then B would have to be the cause of E in case #2 as well; but in case #2, E does not occur. There should be a little slash across the arrow there to indicate that, in this case, E does not

occur. My printer does not do slashes either. This, of course, is the pattern of controlled experiments in science, where you have an experimental case and then a control case, holding all other factors other than the one you're testing constant. And again, if in case #2 here the effect E had occurred, then we would know that A was not a necessary condition, and that we would suspect that B, C, or the combination were the real cause.

Finally, let's look at the Method of Concomitant Variations. This is like the Method of Difference, except that instead of contrasting the presence or absence of A, we vary its magnitude and look for variations in the magnitude of the effect. So, here all the arrows are positive arrows, no slash marks. [end of side A]

[beginning of side B] Now, this is a very powerful method of establishing causality, especially given what we learned yesterday about the relationship between an attribute and its instances. The various instances of the attribute A are related quantitatively, and the attribute is nothing more than the dimension along which they are related; so to understand the causal relationship fully, the causal relationship between A and E, we want to know the systematic relationship between their specific measurements, and the Method of Concomitant Variations is especially well suited to give us that, because we correlate quantitative variations in A with variations in E. More on this point in a moment; I'm going to be coming back to it.

All right, these...this is...I don't want to insult your intelligence. I'm sure you've all either used these methods...I know you've all used these methods, or you couldn't do business in reality. Even simple generalizations rely implicitly on these things. But you probably also learned them in explicit logic or methodology courses. John Stuart Mill, by the way, was not himself an Aristotelian. He was an Empiricist in Hume's tradition. But I think you can see that his methods of induction were a real contribution to the Aristotelian approach to causality. These methods are procedures we use to ensure that our reasoning conforms to the Law of Causality in just the way that the Aristotelian syllogism in deductive logic is a set of procedures for ensuring that our reasoning conforms to the Law of Non-Contradiction.

Now, notice that in all three cases, we had to examine...all three methods...we had to examine more than one case. These cases differ, however, and the differences between them, the presence or absence of A, the variation in the other factors, those differences are crucial to the conclusions that we draw. So, we are not

simply gathering confirming instances, as in Hume's approach. We are looking in a very structured way at a single set of instances.

Now, remember Hume's argument that I gave you earlier, that if there were a necessary connection between cause and effect, we could generalize from a single instance. We wouldn't need to repeat the observation or the experiment. Well, that's true. There is a necessary connection between cause and effect, and since the cases that we compare in using any of the methods constitute a single, connected set of observations, we do in a sense generalize from a single instance. At any rate, we do not rely on repetition. We don't have to collect identical confirming instances.

Now, if I could digress slightly, although I believe this is really just drawing out the broader implication of what I just said: repetition plays no essential role in knowledge at all, not in induction, not in concept-formation, not in any reasoning process. Whenever we reach certain conclusions about a given phenomenon, by observation or by inference, the occurrence of an exact repetition of that phenomenon does not allow us to draw any new conclusions, except the obvious conclusion that this has happened before. Repetition as such, the sheer fact of repetition, is epistemologically barren.

In concept-formation, for example, we could not form the concept "red" by observing two identical shades of red. We form the concept by omitting measurements. But in order to do that, we have to have some measurements to omit. So we have to have shades that differ quantitatively, so that we can grasp the measurement relationship between them. And once we have two different shades that differ quantitatively, and we notice that quantitative relationship, we have all that we need to form the concept. Further instances add nothing essential, although they may facilitate the process psychologically.

And the same is true of induction: observing the same action occurring in two or more identical situations would not allow us to isolate the relevant causal factor. We need some variation among cases in order to use any of Mills' methods. But once we have the kind of variation required by a given method, we have all we need in order to generalize. Further experiments that simply reproduce the cases add nothing essential.

Now, it is true – I hear the question springing to your lips – it is true that the replication of experiments is an important part of science. But that's because replication is a way of checking to see that the

initial experiment was not biased, that it didn't omit any relevant factors in subtle...in all the many subtle ways that can happen. And that is of true, fundamental importance. But the replication does not, strictly speaking, add new evidence. It's only a test to make sure that the evidence provided by the initial experiment was valid in the first place. It's like corroborating one witness's account of a crime with another one. It's a corroboration of the evidence, but not a new piece of evidence.

OK, now, so once we have used Mills' methods to isolate the cause and effect in a single case or a single set of related cases, we can generalize immediately, because the Law of Causality ensures us that the same cause has the same effect. But when we refer to the "same cause," we raise once again the Problem of Universals. In the Method of Agreement, for example, we look for a factor that is common to all the cases, the factor that I've labeled A. A is a concept designating a feature whose instances in the different cases are not exactly alike in specific degree. For example, a college student who tried to explain why he liked certain courses that he took might find that the common factor was that the instructor made the subject matter seem interesting and exciting. That's the cause of his enjoyment. But of course, the instructors do this in different ways, the subject matter differs from one course to another. The student has to abstract the common factor by omitting all those measurements. So, the ability to apply Mills' methods depends crucially on our ability to form concepts, and the validity of our concepts is a crucial part of the process.

So, I want to show, as the final point in our discussion today, the consequences of adopting a bad theory of concepts. I want to look at what effects Nominalism and Realism would have on the field of induction and, thereby, show again one of the virtues of Objectivism.

OK, and let's start with Nominalism, because Nominalism, especially in its extreme form, has an especially destructive effect on induction. As we saw yesterday, extreme Nominalism says that a concept represents an arbitrary grouping of things that have nothing in common, except that we call them by the same name. If that were true, we would have no reason whatever for thinking that everything we arbitrarily group under the concept A would have one of the effects that we arbitrarily group under the concept E.

Here's an example, and, I hasten to add, this is not my example. I doubt I have the imagination to have come up with anything as bizarre as what I'm about to tell you. This is an example put forward by a

philosopher named Nelson Goodman, and it has been discussed at amazing length in the philosophical literature. Suppose we have observed that emeralds are green, and we have observed them under enough conditions, enough varying conditions, that we're sure the color is an effect of the essential nature of the gem. So, we generalize that all emeralds will have the same color. But what counts as the same color? Let me form a new concept, "grue," spelled this way [writes g-r-u-e on the blackboard]. What is grue? To be a member of this color category, an object must be green if it is examined prior to some arbitrary time, let's say 1990; if it is not examined prior to 1990, it's blue. So, this concept designates a color, a set of specific determinate shades of color, which we have separated into two subsets, and those that have been examined prior to 1990 are green, those that have been not examined prior to 1990, which obviously includes all colors that have come into existence after that time, are blue. Now, all the emeralds – some of you are looking at me in amazement; I kid you not, this is a real example – all the emeralds we've examined so far are green. All those emeralds have also been grue; they fit the definition of "grue." If we can generalize that all emeralds in the future will continue to be green, why can't we with equal justification generalize that all emeralds in the future will continue to be grue? But that would mean they will all be blue after 1990, that it's the same evidence, that all emeralds up to this point have been green, supports two generalizations that have conflicting predictions after 1990. Wipe-out! [Laughter.] You could do this, obviously, with any concept.

Now, you might object that "grue" is an invalid concept, because it implies that an emerald has to change color without cause on New Year's Day 1990. That's true, but that's not the fundamental answer. Goodman would reply that the emerald changes color only if we think in terms of "green" and "blue" as our color concepts. If we think in terms of "grue," then an emerald that remained green on New Year's Day would have changed color without cause from grue to non-grue. The fundamental objection to this argument is that it rests on extreme Nominalism. Goodman assumes that any concept is as valid as any other. We can group things together any way we want. We can define predicates, in his language, arbitrarily, without any constraint from reality, any constraint by similarity or any other objective factor.

But on the Objectivist theory, by contrast, as we saw yesterday, concepts must be formed in accordance with the constraint set by the patterns of quantitative and qualitative differences. And the concept "grue" violates that constraint. When we compare two shades of green to something that is blue,

the difference between green and blue is necessarily perceived as qualitative in comparison with the quantitative relationship between the two shades of green. The two shades of green are, therefore, necessarily seen as more similar than either one is to the blue, and that's a perceptual fact, which the mere passage of time cannot affect.

OK, now, if the Nominalist theory of concepts would invalidate induction, what about a Realist theory? Most Aristotelian philosophers who have written about induction, and this includes H. W. B. Joseph, whom I cited earlier, have been moderate Realists, Aristotelian Realists. So, let's take a look at Realism. Now, the major problem that Realism causes in the field of induction is that it isn't true. That's a serious defect, if you're trying to validate a procedure to rely on a premise, if that isn't true. If the only way to show that induction is objective is to bring ourselves to believe in intrinsic essences and attributes that exist intrinsically as abstract, then we have a very tough selling job to do, both to ourselves and to other people, and the fact that, I believe this is one of the reasons why Goodman's argument and other Humean approaches have been so persuasive, that people...most philosophers can't accept that and so feel driven into Nominalism.

However, I want to focus, or turn to somewhat more delimited problems that Realism causes in regard to induction, because I think they bring out once again further virtues of Ayn Rand's theory of concepts. When a Realist says that the same cause has the same effect, the word "same" means "exact identity." The Realist means the possession of the same identical abstract attribute or essence. Suppose we observe that asphalt pavement becomes hot in the summer sun, and after we compare different substances, we isolate the black color of the pavement as at least a contributing cause. So, we generalize that black absorbs more heat from sunlight than white objects. What this means to the Realist is that objects absorb heat in virtue of their blackness, an intrinsically abstract attribute that all black objects share. Two objects may differ in their particular shade or degree of blackness; they may be merely similar. But behind that similarity, behind the concrete, determinate colors we observe, is the abstract attribute that they share, and that abstract attribute, as distinct from the individuating elements that make them individual shades, is the cause of the heat absorption. For the Realist, the fact that the determinate degrees of any attribute are quantitatively related, the fact that they can be ranked in order along a dimension of measurement, is not relevant to the objectivity of concepts or to the process of concept-formation, as we saw yesterday. Nor is it fundamentally relevant for the Realist to causality or induction.

Now, one consequence is that Realists view abstractions as the real causal agents in nature. An entity does not act in virtue of its specific determinate nature; it acts as an embodiment of an abstract attribute that it shares with other entities. And another consequence is that Realists gravitate toward a qualitative understanding of causality. I think this is most easily visible in the science of the period prior to Galileo, and Galileo's insight that language is the mathematics of...is the language of nature. I'm not sure what I just said. Mathematics is the language of nature, was Galileo's insight. The Aristotelian scientists, and this, by the way, is not...I'm not charging them with not grasping something that they really weren't in a position to grasp yet, but the fact is, they typically sought to establish qualitative laws, such as: heavy objects fall toward earth, species reproduce after their kind, fire burns, water flows. The specific quantitative measurements of an attribute did not play any major role in the theories. But those specific measurements are crucial. Massive objects fall to earth as a result of a specific degree of gravitational force that is mathematically related to the mass of the object. In general, scientists look for the causal relationships which show not just one variable or one attribute, X, depends on Y – that's the qualitative causal relationship. They look for functional relationships [writes equation on blackboard] of that form, which relate the two variables X and Y in a way that reveals the quantitative relationship, the systematic, mathematical relationship between the particular values of each variable, the specific measurements. You know, for example, the Law of Gravitation [writes equation on blackboard], if I remember my physics correctly, looks something like that, and the gravitational attraction between two objects is related by this function to the masses of the objects and the distance between them. Or, in economics, some of you have heard Nort Beuchner's lectures on the equation of exchange, which is that money supply times the velocity equals price level times the number of transactions. OK, again, and I'm sure there's a constant in here somewhere, again, what we're doing is saying there is a systematic relationship between the specific measurements of those attributes.

Now, the Objectivist theory of concepts is ideally suited to explain this aspect of scientific inquiry. We form concepts by noticing quantitative relationships among similar objects, relationships that allow us to rank them on a common dimension, and measure them by a common unit. Induction, then, is a matter of correlating two such dimensions. We don't need to attribute causal efficacy to abstract attributes that lurk behind the specific measurements. And yet, the classification of similar objects and the omission of their

measurements is governed by the objective existence of their quantitative relationships, so it is not an arbitrary process, as it is in the Nominalist theory. And we are, therefore, immune from the skeptical conclusions that Nominalists are forced to draw about causality and induction.

OK, now there are many other questions about induction that I have not addressed today. There are questions of detail, about the application of Mills' methods, questions about how you know when you have varied all the factors, and so forth. There are also questions about how we formulate and test theories that are posited to explain causal connections. These are all important questions in the philosophy of science, but the point of departure for any proper philosophy of science, in my view, is the Aristotelian view of causality and Ayn Rand's theory of concepts. Together, these provide a validation of the Law of Causality and of the basic method of induction, which are the foundations for everything else. Thank you. [Applause; end of lecture]

### Question Period

Thank you very much. I think we have almost as much time as I was supposed to leave you for questions, and I'm sure there are some, so fire away. Who wants to start?

Yes. [Male questioner: *I've got one little one and a bigger one. When you say that the Hume people would say, the grue/non-grue type of thing in 1990. I understand what I'd say to that, but I didn't understand what their...when you said that.*] OK, what would you say? [Well, I'd say, you're asking if...the laws of reality you're going to say...remain green the entire time, and you're almost playing with semantics. I wouldn't say that...now, that's the essence of what it'd be thinking. It's going to be green; it is going to be what it is, and it won't change in 1990, so it's just a matter of what are you going to call it...] OK, Goodman's claim, in answer to that, is that, as I said, you think in terms of "green" and "blue," that's your parochial, Western assumption. I think in terms of "grue." In order to remain grue, however, a thing has to act a little differently than it does if it has to remain green. If it stays green, in your terms, that means it changes color in my terms, from grue to not...actually he defined "bleen" as the exact opposite of "grue." [Laughter.] So, if an object stays green past that point t, 1990, then in Gibson's [he means Goodman's] framework, it changes from grue to bleen. Now, I don't want to defend him too far. Clearly, where we're talking about

a single object with a single specific color, if you were to observe that continuously from prior to 1990, you know, over New Year's Eve of 1990, OK, one of two things would...well, a number of things might happen. We know what *would* happen. It would stay the same color. But Gibson, or Goodman would say we could observe then whether it remains grue or remains green. Now, it does seem to me that the whole argument presupposes that we can tell the difference between blue and green. Otherwise, there's no difference between green and grue after 1990. So, given that we can tell the difference on the stroke of midnight, 1990, we will be able to observe, we will be able to discriminate the cases, of it remaining green and it remaining grue. And we will be able to compare the color that it is right after midnight with the color that it is just prior. So that Goodman himself, by assuming there is an observable difference between these two colors, particularly at the level of determinate colors, is committed to our being able to tell...it...his...it implies that there is an objective difference between it staying green and staying grue, because since we can tell the difference between green and blue, we can tell if it changed from one to the other at that point. That's a real change, whereas remaining green is not a change. So, in that sense, there's an asymmetry between the two cases. But more generally, the problem here that grue introduces is not concerned so much with the Law of Causality applied to specific, determinate measurements, but rather to types of measurements. It's...the problem: how can we generalize from this specific shade of green to all shades of green? Or from this specific mass to all measurements of the variable "mass." Clearly, in doing that, we are forming a concept, and we are putting together things that differ, so now we're dealing with two shades of green that differ. When we say that all emeralds are green, we are treating, lumping together two shades that in fact differ at the level of determinate, concrete shades. What justifies in putting those two together, as opposed to the blue and the green? And there the key question is: is there an objective similarity basis for saying the two shades of green are more similar than the two shades of, than the blue and the green? Goodman, as an extreme Nominalist, says no, and the Objectivist theory says yes. Goodman, by the way, believes that any two things are exactly as similar to one another as any other two things. You know why? Here's the argument. Take any two things, and then take the entire universe of things which, let's say, is finite. Things can, all the things in the universe can be grouped into sets, and there, you know, by some mathematical law involving the exclamation mark, we can know the total number of sets it is possible to construct out of all of the objects in the universe. Now, two objects are similar if they can be described by the same term. Similarity, says Goodman, is relative to our



conceptual framework. And a concept is just an arbitrarily collected set of objects, so any two objects are members...two objects are similar if they're members of the same set, and any two objects are members of exactly the same number of sets, so they're exactly...any two pairs of objects are exactly similar. My watch and your watch are as similar as my watch and the Empire State Building. That's Primacy of Consciousness with a vengeance.

OK, yeah, let me take another one, another question, yes. [Male questioner: *Hume says that axioms were arbitrary, but he still didn't accept axioms. I have trouble understanding that. I would say to Hume, well, if axioms are arbitrary, then if you don't assume A is A, I don't understand a word you're saying. I mean, you can't speak the language without assuming A is A.*] Well, Hume would probably say, fine, that's true. But that only...that's a fact about language, it's not about reality. I mean, for Hume and that whole approach, axioms are conventional, and they may be necessary to the practice of speaking a language, but after all, that's a human activity which we construct on our own. I mean, suppose...they are like the rules of a game, to use Wittgenstein's way of putting it. Suppose you slid into second base, and the second baseman was waiting there with the baseball and tagged you, and the umpire said "You're out," and you said, "Wait a minute! Prove to me that being tagged before you hit the base objectively makes me out. What's the basis for that? How can you validate this rule?" Well, that would be an absurd question. It's just part of...that's part of the game. It's arbitrary, but if you don't want to observe it, then don't play the game. Hume and that tradition says the Law of Identity is one of the ways we organize our language, true enough, and we couldn't speak without it. But...we couldn't. Maybe someone...some other language with a different set of rules could. But that only tells us something about our language, not about reality. Now, you're right, though, to push the question, because a key part of saying that this is conventional is that we could...there could be another language structured around a different set of principles, and in fact that really can't be done. Unlike baseball, where you can imagine other games, you can't imagine...with different sets of rules, or you can imagine a similar game like baseball only you get to second base free if you're...it's a wild base. But we can't imagine a wild language, where the Law of Identity doesn't apply. And it's not just a fact of imagination. I mean, you can show that even someone who tries to assert the possibility of such a language has to assume the Law of Causality and thereby reaffirm it in the very act of denial.

Yes. [Male questioner: *Does your point about the...of nature of induction ascribe a lesser importance to induction in terms of gaining knowledge?*] The point about repetition? [Right, that you don't gain any new knowledge by new instances of some evidence.] No, I'm not saying...no. Repetition, I'm saying, repetition as such is not a part of the inductive process. The inductive process is a crucial...it's a foundation of all our knowledge, all of our general knowledge. But the point is simply that we don't rely on repetition as such, when we use that process. So, I'm not impugning the importance of induction. I'm just saying that it doesn't involve any reliance on repetition.

Let me see if I can get another...yes. [Male questioner: *...a little bit of a problem with the phrase "it is in the nature of an entity to do something under certain conditions." I've always felt it's in the nature of the laws of the universe for an entity to act a certain way under certain conditions. For example, if I throw anything, it will continue forward because of the Law of Momentum, so it is not in the nature of a ball to carry forward; it is in the nature of the Law of Momentum to carry any object forward.*] Yes, right, and this is a good point, and it brings out once again the relationship between concepts and induction in the following way. It is not in virtue of being a ball that an object continues in motion unless a force acts on it, and we can prove that by the use of Mills' Method of Difference. We take something that is not a ball, and it continues to move forward, it has the same effect. It continues to move forward if there's no interfering force. So, it's not a rounded, rubber object that it has that, and yet it has to be some aspect of reality. That is, the law is not something that's inscribed on our minds or on the skies that we read and that, and which things obey in the way that we obey legal laws out of fear of the consequences. Things obey laws because their natures require them to. In this case, and it is, in this case, it is an aspect of a thing's nature, namely, its mass which happens to be something that all physical objects have in common. So, we have to reach the appropriate level of generality in order to identify the relevant causal factor in the thing's nature. We have to reach the relevant level of abstraction, which in this case is as high as you can go in science, anyway, so far as we know. But it's still an aspect of a thing's nature. It has to be existence, in some form, that makes things act the way they do.

Yes, Patricia. [Patricia: (inaudible).] The question is, determining what factors are relevant, and the example that Patricia mentioned was the generalization: All swans are white, which has always seemed to be a shaky generalization, even though we have thousands of, you know, countless confirming instances of it, with all kinds of variations of species and circumstance and

so forth. And this is one of the kinds of problems that I alluded to at the end of my talk as something that a philosophy of science has to provide some specific guidelines on. There is tremendous variation in the amount of structured inquiry structured by Mills' methods that you need to undertake to support different kinds of generalizations. The classic opposition is between finding the melting point of a metal, where all you have to do is melt one sample, and you know that any metal will melt at that temperature again, and this particular exact case of "all swans are white," where, you know, there was abundant evidence, and yet it finally...it was...always seemed shaky, and it eventually did turn out that there are black swans, you know, I believe in Australia. I hope that's where they were; that's what I said in my book. [Laughter.] But I think the key factor here is that in the case of the metal and the melting point, we have a good understanding of the mechanism by which a metal melts, and that understanding of the mechanism shows us that it's in the atomic structure of it, the thing that defines it as that kind of metal, that governs the way it changes state, so all we have to do is find what the melting point is for one particular instance, and we can generalize immediately. That's because we have an enormous context of prior established theories about the melting points of metals. Prior to that point, for example, in the...before the elements had been isolated, where you dealt with metals in different ores, well, those different ores do melt and behave differently at different temperatures, so you would have had to do a lot more varying of factors. [*Plus, thermometers didn't exist back then!*] A lot of varying of factors was necessary in order to extract the – conceptually as well as technologically – the metal from the ore. In the case of white swans, it's not clear what the biological basis for the coloration is, or it's not clear why all species that...all birds that are members of this species, using the standard characteristics that seem essential to species membership here, how those relate to the coloration. So, we are left in the position of having...knowing that there are...precisely because we don't know the mechanism by which these things are related, we don't know...we know...we do know that we haven't identified yet all the factors, and so we have to make an explicit proviso on that generalization that this is subject to our current context of knowledge. That's the best answer I can give you at this point. OK, thank you. [Applause.]