Objectivist Epistemology in Outline

Gregory Salmieri, July 2006
Lightly edited tape transcript
Last revised: 2/11/2010

Lecture 1: Prelininaries

1.1. Introduction

1.1.1. AR on the nature and importance of epistemology

Ayn Rand defined epistemology as, "a science devoted to the discovery of the proper methods of acquiring and validating knowledge." She said that "epistemology is the foundation of philosophy", and that it is with "a new approach to epistemology" that "the rebirth of philosophy" has to begin. In notes written to herself, she wrote that philosophy is primarily epistemology.

And yet, Ayn Rand's epistemology is comparatively unknown. In the popular culture, she's associated with her political and ethical views. Objectivists usually know something about her epistemological ideas, but even amongst Objectivists—people who use Objectivism as a guide to their lives—, I don't think her epistemology is understood in the way her ethics is. My goal in this class is to give an overview of Ayn Rand's new approach to epistemology, to foster a greater appreciation of her achievement, and to provide a context and orientation, or a greater context and orientation for people who want to study and think about it further.

1.1.2. The nature of ITOE and other sources

Let's start by saying something about the nature of the sources—that is, the works in which Ayn Rand commented on epistemological issues. The major work on epistemology is the *Introduction to Objectivist Epistemology* (ITOE as we call it). This is neither a systematic treatment of epistemology--it's title suggests that it's not--nor is it, as the title *Introduction* might suggest, a survey of the essential points in her epistemology, an overview of it. No, rather ITOE presents a "summary" of one "cardinal element" of her epistemology, namely her theory of concepts—and it presents a summary of this cardinal element "for the benefit of philosophy students." What she does in this work is introduce her distinctive epistomology to *students of philosophy*—people who are already acquainted with general philosophical issues, and with the nature of the field, people who are in a position to already know what the problems of the field are and, therefore, to see how the theory she presents in that work, a new theory of concepts, provides the key to the field—how it is the beginning of a new direction in epistemology. She indicates in that work how her theory of concepts is fundamental to other issues in philosophy, but these are indications—things she mentions in the course of laying out the theory. She doesn't give a systematic account of what the field would look like as a whole. She doesn't show how the theory of concepts underwrites solutions to the various problems that epistemologists have faced, and she doesn't show how the solutions to these different problems cohere into a system—a new approach to the discipline as a whole.

These are tasks she left to her students. Both of them have been advanced incredibly by Dr. Peikoff. The first task—that is, the task of applying her theory of concepts to different issues in epistemology, he took up in, for example, "The Analytic-Synthetic Dichotomy" article, which Ayn Rand published as an appendix to ITOE, and, of course, recently, on a larger scale with his lectures on induction where he talks about how the theory of concepts leads to a new approach to induction. And,

of course, in his work *Objectivism: the Philosophy of Ayn Rand* (OPAR), Dr. Peikoff gives us a systematic presentation of the philosophy of Objectivism as a whole, including its epistemology. I'm going to draw, as we'll see, quite considerably on the systematization that Dr. Peikoff provides in that work. One other notable work by a student of Ayn Rand's that I want to mention because I'm going to rely on it heavily here (and becauuse I want to recommend it to people) is Dr. Binswanger's taped lecture course *Consciousness as Identification*, which brings out a theme in Ayn Rand's epistemology that I'm going to be talking about today and throughout the course, and I'll mention that theme when I get to it.

1.1.3. The nature of this course

A few points about this course: First, I'm going to be less conservative, in particular than Dr. Peikoff was in OPAR. That is, I'm going to include a lot of material that I don't know for sure that Ayn Rand would have agreed with. For example, ideas of Dr. Peikoff's that he says he didn't include in OPAR for this reason, also some of his more recent ideas and ideas of my own. I think all these ideas help us to see the power of the Ayn Rand's theory of concepts because they're applications of it, but they're not parts of Objectivism proper. Second, I'm going to be focused more on the connections between different points than on proving each point individually. Again, what I want to provide is an outline, a kind of context: an overview of a whole that one can hold in mind when one then goes and studies parts of this whole in greater detail.

1.2. Consciousness

The first substantive section on the outline after the introduction, and the one I'm going to turn to now, is the topic of consciousness in general. I think the widest thing we can say about Ayn Rand's approach to epistemology is contained in a quote from the last chapter of the *Introduction to Objectivist Epistemology*: that the fundamental principle here is that, "Existence to be apprehended must be obeyed"— "The rules of cognition must be derived from the nature of existence and the nature, the identity, of his cognitive faculty." The rules of cognition must be based on the nature of existence, and the nature of consciousness. What then *is* the nature of consciousness?

Well, ITOE proper (that is after its preface) begins, "Consciousness, as a state of awareness, is not a passive state, but an active process…" And she goes on to describe the process, but we'll stop here for a moment. This is the point that Dr. Binswanger was instrumental in directing at least my attention to, and I think a lot of other people's. "Consciousness, as a state of awareness, is not a passive state, but an active process." Now to understand this we have to break it down, and I want to start by talking about what it means to talk about consciousness as a state of awareness.

1.2.2. Consciousness as a state of awareness and as a faculty.

We can distinguish two related senses of the word "consciousness": there's consciousness as a state of awareness, and consciousness as a faculty. Now "consciousness" in either sense is axiomatic. That means that consciousness is directly known and that it cannot be defined. So I can't explain to you what it is; all I can do is direct your attention to the relevant facts that fall under this concept of "consciousness".

So what is referred to by "consciousness as a state of awareness"? Seeing—what you now see, you're now seeing, and your seeing is a state of awareness. Hearing, and in general *perceiving*, are all states of awareness. In addition to these, there's awareness in *conceptual* form, in the form of concepts

ordered into propositions. We'll discuss this form of knowledge in great detail later, but for now we can just think of it as the kind of awareness we have—the kind of knowledge we have—when we think or talk in sentences. So not only do you see me, you're aware that I'm a person, that I'm a philosopher, that what I'm doing is lecturing, and so forth. And that is awareness. The awareness you have is consciousness. Synonyms for this sense of "consciousness"—consciousness as a state of awareness—include "awareness" and also sometimes "perception"—using the word "perception" in a wider sense than the sense in which we use it to refer to sense-perception (i.e., to the information you are aware of through your five senses).

We need to contrast consciousness as a state of awareness with consciousness as a *faculty*. In Galt's speech, Rand says that there are two corrilary axioms implied by the grasp of the fact that existence exsits. One is: "…that one exists possessing consciousness, consciousness being the faculty of perceiving that which exists." Now perceiving in this sense is used in the widest sense to mean "being aware of". So the second sense of "consciousness" is the faculty of awareness. "Consciousness" means either awareness or the *faculty* of awareness. Now a "faculty" (this is my definition) is an enduring attribute of an entity (especially an organism) in virtue of which it is able to engage in some activity or set of activities. Your faculty of consciousness is an enduring attribute of you, something you have even in those moments when you're unconscious, and in virtue of it you are able to engage in some activity or some set of activities. Other examples of faculties include the nutritive faculty which all living things have—the ability to process nutriment into more of themselves—the reproductive faculty, the faculty of locomotion, and so forth. We're focused, though, on the faculty of awareness.

1.2.3. Awareness

1.2.3.1 Awareness as an active process

Now consciousness in the sense of a state of awareness is an active process, it's an *exercise* of your *faculty* of awareness. Consciousness in the sense of awareness is not primarily something one *has*, it's something one *does*. Seeing, for example, is a complex action consisting in all sorts of processes, primarily performed by the brain, and the same holds true for each of the other senses. So consciousness as a state of awareness is an active process, it's something we do. Described holistically, it's the activity of grasping the identities of things. Thus Ayn Rand's famous line "consciousness is identification."

Now "grasp" here is a metaphor. We can't define awareness, but we can put the point in different ways and draw attention to different aspects of it. Consciousness as a state of awareness is an active process and what it is is a grasp of the identities of things. The metaphor of "grasp" is very fruitful. I've now grasped this cup, but my grasp of it isn't just my initial action of taking hold of the cup. As long as I hold the cup in hand, I'm continuing to grasp it—there's an action that I'm continually doing: I'm exerting pressure on the cup in a certain way, my muscles are at work, etc. The grasp is an ongoing process. We speak of "having a grasp" on something, and that's fine, just like we speak of "having an awareness" of something. And we should speak like that, but we have to keep in mind that there's a process of sustained action here. The grasp is something that we do, and likewise with a metaphorical "conscious grasp"—with an awareness of something.

1.2.3.2 Awareness as having identity (the form/object distinction)

So, awareness is an action. Moreover, each act of being aware is a specific act, which the faculty of awareness is able to perform because of its nature. We can see this again by analogy to physical grasp. There are different forms of physical grasp. Right now, I'm holding this cup in what

people who study the hand call a power grasp. Now I'm holding it in what's called a precision grasp. And there are tons of different forms of grasp, which are subdivisions of these two basic forms of grasp. Likewise, there are different *forms* of being aware of something. Seeing something is different from hearing it, and in general perceiving something is different from knowing it through concepts—for example, inducing or deducing it. These are different forms of being aware. Moreover, each act of seeing is different from each other act. It has a different vantage point, a different level of clarity, etc.

I've been using the concept "form" here, and it's important to distinguish the *form* of an awareness from its *object*. A form of awareness is the identity of a state or act of awareness (states and acts are the same because states of awareness are acts) as distinct from the identity of *object* of which one is aware. When you're aware, you're always aware of an object; the form of your awareness is the identity of the state of awareness you have as distinguished from the identity of its object. To use other terms that are often used for this distinction in Objectivism, the form of awareness is the "how"—it's how you're aware. The object of the awareness is the "what"—what you're aware of. You can know, for example, the same fact in two different forms. You can know visually and tactily the shape of an object.

Now, form and object are constantly confused or conflated in philosophy, leading to all kinds of problems. They're turned into one package-deal: object-in-a-form (think of it with hyphens between it). For example, when I see a man blurrily because my vision is impaired, to conflate the form and the object would be to think that I see a blurry-man.

Now if you make this conflation—if you package-deal the form and the object—there are two things you can do with this package-deal: you can either think of it as a feature of consciousness or as a feature of existence. If you think of it as a feature of existence, at least at the perceptual level, you're making an error called "naive realism". (Of course, it's not called an error by people who hold it, but it is an error.) It's the view that consciousness is aware "no-how", in no form, of an object. But the object—what is thought to be the object—already has the identity of the form smuggled into it. So you think, for example, that there's a blurry man somewhere who I'm aware of when I take off my glasses.

The alternative view is attributing the package-deal (the blurry man) to consciousness, and this is subjectivism (particularly at the perceptual level, which we're using as an example now, "perceptual subjectivism"). It's the idea that there's a blurry man in my head, and that's what I see, and my awareness of this blurry man in my head is not an awareness of anything in the world. According to this view, I'm aware "blurrily" and "manily" of nothing.

There are other views that one will find in the history of philosophy, but they're all halfway houses of some kind or another between these two views (or almost all, I'm speaking very abstractly now about different views in the history). They're all halfway houses—things that you get to by combining elements of each.

The correct view, Ayn Rand's view, is that consciousness is awareness of *something*, *somehow*: of an *object* in a *form*.

(I want to note parenthetically—and people can ask me in the question period—that throughout the history of philosophy, philosophers are constantly starting to grasp this distinction, and then for certain reasons not being able to hold onto it, so they lapse back into one of these two errors. And philosophy is filled with the positing of existents like Platonic "forms", Fregean "senses", "meanings" and so forth. If you're familiar with the history of philosophy, you can see how these are all confused attempts to distinguish form and object which then lapsed back into making forms into extra objects. But since this is a historical, comparative point, I don't want to discuss it here at length. People can raise it in the question period if they want.)

1.2.3.3. Awareness as volitional at the conceptual level

So far we've been talking about consciousness. Consciousness as a state of awareness, we've said, is an active state with identity; it's an activity, and that activity has an identity—it has a form. Now we want to add that, at the conceptual level, awareness in terms of concepts is *volitional*: it is performed consciously and by choice as opposed to perceptual awareness, which is not performed consciously and is performed automatically. For example, if I held up a first edition of *Atlas Shrugged* in front of a child (or in front of anybody who didn't already know a lot about it), he would automatically see it. He would not automatically however know 1) that it's a first edition, 2) that it's valuable 3) that it's a great work of literature, or 4) that it presents a true philosophy. In order to be aware of any of these things, he would have to initiate, sustain, and direct *by choice* a process of being aware, and that's the process that we're going to study in this class: the process of conceptual cognition. Actually we're going to study it in the next three days of this class; today we're going to talk about some things that we need to understand in order to be in a position to talk about it.

1.2.4. The need for epistemology

Because we need to initiate, sustain, and direct conceptual awareness by choice, we have to know *how* to be conceptually aware. We need a *method*. It's something we do, we have to direct ourselves in doing it, therefore we have to know how it is that we do it. And this is the job of epistemology.

Now, of course, during childhood, everyone acquires some implicit knowledge of how to know. And everyone uses this knowledge to discovers and holds some knowledge of the world without explicitly identifying how he's doing this. This implicit method, however, is not sufficient. It's not sufficient for a human being for two reasons: because human beings are neither omniscient nor infallible. This is a point Ayn Rand makes always in the context of discussing the need for epistemology. And let's break it down.

What does it mean to say that man is not omniscient? Well, often we just find ourselves knowing things. For example, if I held up a copy of *Atlas Shrugged,* all of you would recognize it as a copy of *Atlas Shrugged*. It's not that it's like your vision of it that happens physiologically (we'll talk about the fact that vision is physiological in a moment), but nevertheless the knowledge would just come to you, you wouldn't have to work for it; you would just, in effect, find yourself knowing it. But there are many things that we don't just find ourselves knowing. We don't, for example, just find ourselves knowing how to get to the moon, what diet is the most healthy, what political candidate is the best, or whether the dangers posed by Christian fundamentalism are greater than those posed by Islamic terrorism. These are things that we don't just find ourselves knowing, so we need to find out how we can find the answers to these questions because we're not omniscient.

Moreover, we are not infallible. This means that we can make mistakes, so that some of the things we just find ourselves "knowing" might not be genuine knowledge. Some of us, when we pole our subconsciouses, might just find that we "know" that, say low-carb diets are the healthiest way to lose weight. But that might not be actual knowledge. Because we're not infallible, we need a method to validate our knowledge, not just to discover it.

Notice that arguments about, for example, which diets are most effective, which religionists are more dangerous, and so forth, at least when conducted amongst knowledgeable people, often turn as much on issues of *how* to prove—*how* to validate—the various conclusions that are at issue as they do on the actual evidence involved. For example, on HBL, the email list, there's been a sustained discussion about the comparative dangers and evil of different religions, and what people are disputing is not the content of the various sacred texts, but to what extent the context of the texts *matter*—what kind of inferences we're entirled to draw from them, etc.

So we need epistemology—we need an explicit method—, because we're not omniscient and

we're not infallible. To package this point with a colloquialism I don't really like: we need epistemology because of what we don't know and because of what we "know that isn't so".

1.2.5. Definitions of "cognition" and "knowledge"

Now of course, and this brings us to the need to introduce some new concepts, "know that isn't so" has to be put in scare quotes, because you can't *know* something that's not so. We need a concept wider than "knowledge" that includes knowledge (i.e., awareness) and also errors; we need a concept for *attempted awareness*; and that concept is "cognition". That's the word I'm going to use in any event. Though it's a bit awkward, I'll speak of "a cognition", for example, in the way that I speak of "an awareness". Cognition is attempted awareness. The concept subsumes both awareness and failed attempts at it.

At this point I also want to say a little bit about the definition of the concept "knowledge". You'll notice that I've been sliding back and forth very freely between the words knowledge and awareness. I think there are three senses of "knowledge". They're all perfectly valid, they're closely related, and all of them are used by Ayn Rand at one point or another in her writing. I think the widest sense of "knowledge" is simply synonymous with "awareness"—the axiom of consciousness. Then there's a narrower sense in which "knowledge" refers only to stable and enduring awarenesses, as opposed to, for example, a fleeting glimpse of something out of the corner of your eye. And finally, the narrowest sense of "knowledge" refers just to propositional knowledge: knowledge held in the form of sentences made of up words which words denote concepts. I think it's usually clear from the context which is meant. I will sometimes use the word in each of these three senses—Ayn Rand does. If people want examples of her using it in different senses, you can ask during the question period.

1.2.6. The faculty of consciousness

So we've spoken about consciousness as a state of awareness. I want to talk a little bit about the faculty of consciousness now because there are some things we need to be clear on about it, otherwise they might result in confusions later.

First off, the purpose of consciousness is not just being aware of the world as an end in itself. The purpose of consciousness is awareness of the world to direct action. And as a result of that, many of the states or actions of the faculty of consciousness are not themselves cognitions—they are not themselves attempts at being aware of anything. All states of consciousness have a content which is derived from cognition, but not all states of consciousness are attempts at being aware of an object.

Here are some states or activities of the faculty of consciousness that are not cognitions: imagining, desiring, initiating an action, dreaming. Each of these actions has a content. In all cases, the content has to have derived from prior cognition: you cannot imagine something if you don't know anything—if you've never perceived anything—nor can you desire something if you've never perceived anything, nor can you initiate any actions. Nevertheless, when you imagine something, the act of imagining it is not an act of discovering something new, or an act of coming to have greater knowledge in any way.

The states of a consciousness that are not cognitions are all involved in some way or another with the use of cognition to guide action: they are either means of using cognition to guide action, or they're a byproduct, somehow, of cognition or of action guidance. (Dreams presumably are some kind of byproduct.)

Okay, so we have many states of consciousness, not all of them cognitions. All of them have some content. Consciousness as a whole, I think, can be divided into levels depending on the type of conscious content that is used in any given state of consciousness. And there are basically two levels

(there's a wrinkle on this which we'll get to later, but there are basically two levels): *the perceptual level* and *the conceptual level*, depending on whether the content involved in a conscious state is perceptually held content (content in a perceptual form) or conceptual content (content in a conceptual form). For example, Animals percieve, and then they perform a number of other mental actions with their perceptual content. They remember it, they associate perceptual memories with things they're now perceiving, they desire things in perceptual form, and I think, in a limited way that we'll discuss later, they can even project the future—the very near future—in a perceptual form and initiate action. Human beings do all these things in conceptual form. If you want an illustration of the difference between conceptually and perceptually desiring something, think of desiring a piece of cake, which you can do by imagining the crumbliness and the texture, you know, literally imagining it and imagining what it would taste like, or you can do it by thinking "cake". There are certain things, of course, that you can only desire conceptually, like, say, the triumph of Objectivism in the culture—there's no image of that that you can have. You might associate some image with it, but you have to have the concept to have the desire. This is not the case, however, with desiring a class of water when you're thirsty.

1.3. Perception

Okay, I want to turn now to the topic that's going to occupy our attention for the bulk of the rest of the day, and that is the topic of perception. I'm going to cover it a little more hastily than I might otherwise have, because we've had a lecture on perception already at this conference, but I think I need to address it in particular from the perspective of what's going to follow (and also people are going to buy this on tape and they won't have just heard Dr. Binswanger's lecture on perception).

1.3.1 The relivance of perception to epistemology

What is the relevance of perception to epistemology? That is, why are we discussing it here? Well, Ayn Rand says, in the *Introduction to Objectivist Epistemology*, page five, that "…epistemologically, the base of all of man's knowledge is the perceptual stage." That is, it's perception. Perception is what all our other knowledge is going to be grounded on. And so, the basic advice that epistemology gives you about how to conduct conceptual level cognition is to base it on your perceptual knowledge. We're going to see what that means and how we do that going forward. But in order for that to have content, we need to know two things: one, we need to differentiate perception from some other states that it might be confused with, so that we understand what it is that epistemology tells us to base our knowledge on, and we also need to understand why perception can serve as a foundation for knowledge—what it is about perception that enables it to be the thing on which we base all our other knowledge. We'll find that, if you confuse perception with any one of a number of different things, you will come to think that perception either is impossible to build on, or else that it can be built on, but what you'll be building on is a bunch of errors. So what we're going to do is distinguish perception from several different things. From sensation, from perceptual judgment, and from something that I'm calling "visualization" which we'll get to towards the end of today. And we'll draw some conclusions out from these differentiations.

1.3.2. Perception as distinct from sensation

What, then, is a sensation? Ayn Rand says the following in *The Virtue of Selfishness* in "The Objectivist Ethics": "A sensation is produced by the automatic reaction of a sense organ to a stimulus from the outside world; it lasts for the duration of the immediate moment, as long as the stimulus lasts and no longer."

To get clear on what a sensation is, let's imagine the state of mind of a creature that, plausibly (I'm not sure if this is true), experiences only sensations. Think of a scallop—you know, the little shellfish the muscle of which we eat. A scallop, let's say, only experiences sensations. We can think of the scallop as a big tongue, just exposed to the world—a tongue not connected to any brain that synthesizes what it gets with everything else, the scallop is just a tongue, and the scallop's awareness consists of the taste of where the scallop is. So we can think of the scallop's mental life as consisting of an alternative between "yum" and "ick", which in each moment is just the current "yum" and the current "ick". The scallop can use that big muscle to open and contract its shell and scuttle along the sea floor. It can't experience where it's going when it does this, just whether it's yummy or icky where it is at any given moment. If it's in a yummy place it stays, and if it's in an icky place it scuttles along. (At least that's how I imagine scallops function. It doesn't matter if this is true. The point is just to give you a sense of what it would be to live on the sensational level awareness.)

The scallop's "yum" or "ick" sensation is a primitive form of awareness: it's a differentiation between two states—a sort of low-level knowledge of the world. (And incidentally, Ayn Rand once described sensations as a form of knowledge.) And that's all the scallop has. Now contrast that with the knowledge of a seagull, for example, which swoops down and eats creatures like scallops that are functioning in this way. While the seagull is flying over the water, it sees the water below it, by looking at the water over time; it identifies a creature in moving through the water; then it guides itself towards that creature, tracking it as it moves; it grabs it, picks it up out of the water, and drops it on rocky ground to break its shell by the impact; it tracks it as it falls down, then it flies down to where the thing was, picks it up, eats it, feels the thing going down its throat. Thoughout this sequence, the seagull has a constant, continuous experience of the same object over time passing through different stages as the seagull interacts with it. This is the difference between sensation and perception.

"A 'perception'," says Ayn Rand in the same place in *Virtue of Selfishness*, "is a group of sensations automatically retained and integrated by the brain of a living organism, which gives it the ability to be aware, not of a single stimuli, but of *entities*, of things." There are two main points that we need to get about perception as opposed to sensation: one, that a perception is an integration, not just a differentiation, and, two, that it's performed by the brain: it's a physiological integration. First, the point that it's an integration.

We now can finish the first sentence of Introduction to Objectivist Epistemology. "Consciousness," we have seen her say, "as a state of awareness, is not a passive state, but an active process that consists of two essentials: differentiation and integration." Differentiation is telling things apart, as the scallop does when it tells apart the yummy places from the icky ones. Perception is the integration of sensations into a form of awareness in which we can differentiate *entities* from one another. Integration is the putting of things together into a new whole, not just a collection or a heap, but something that has coherence to it, that's a "one". Perception is an integrated form of awareness. It integrates sensation, or at least data coming from our different sense organs, into a holistic awareness of entities as distinguished from one another and from their background.

This is as opposed to a common false view of perception which thinks of it as a whole lot of sensations exerpeinced at once (or in succession); think of the scallop's "yum" or "ick", but along with it you get "red" or "white", "loud" or "soft", you know, a whole bunch of these alternatives coming at you, and then you've got to put them together somehow. Perception is the integrated form of awareness—the integration of all these sensations, and it occurs physiologically—it's performed within the brain. To see the importance of this point, suppose you were just getting what you got in consciousness as you looked out at the world was all these little atomic sensations of the form that the scallop got, you know: "ick", "gray", and so forth, in each of your sense modalities. You would have to somehow, if you wanted to be aware of a world of existents (of a world of entities outside of you), you would somehow have to process all of those sensations, you'd have to somehow infer from that vast

amount of sensations to a world out there, to the entities that are sending you these signals.

But it would be impossible to do that without conceptual knowledge, and people who have this view of perception—that is the view that there is no such thing as perception as distinct from the reception of sensations—inevitably come to the conclusion that we either have innate concepts which are somehow at work in figuring out that there are entities out there. The concepts must be innate, because we couldn't get these concepts out of sensations, since we would need the concepts to get anything out of sensations. Therefore, if you conflate perception and sensation, you lose the ability to see perception as the base of knowledge. Either you become a skeptic and you don't think we can know anything (*a la* Hume), or else you think we can, but only because we have automatic concepts somehow built in to us (*a la* Descartes, or in another way Kant—although at a deeper level he's a skeptic). So it's very important to understand how perception can be the base of knowledge, that it is distinct from sensation, and that it is a physiologically integrated form of awareness, a form by which we're aware of entities, and that this happens automatically; it's done by the brain.

Now what's the proof that it happens automatically and it's done by the brain? Well, you have to learn that there are sensations. You don't automatically isolate red and blue and loud and icky, you're aware of the world as a whole. It takes work to be able to to isolate these sensations, to think about them individually. It's something that only adults can do. Little kids have to learn about the objects around them, and only much later can they come to think about these individual sensations. In fact, it's an inference that anything is aware of individual sensations. God knows what a scallop's awareness is like. I mean, I don't know—I don't even know how it can be figured out, but if it can be figured out, it can be figured out only by scientists studying scallops in some complicated way. In this connection, incidentally, Ayn Rand thought that infants go through a stage when they're experiencing sensations only, before their brains become, in effect, calibrated to integrate sensations into percepts. She thought this based on the psychological evidence of her day; current evidence indicates otherwise; at least as I understand it, it indicates that infants are from the beginning at the perceptual stage, that they're perceiving whole entities and tracking them as they move.

There's one more point before we move on to distinguishing perception from perceptual judgments that I want to make under the broad heading of distinguishing perception from sensation. Sensations are isolated in a moment, perceptions are awarenesses of entities which are things that exist and move over time. And it's under this heading that we can see the point that Dr. Peikoff's made in his lectures from on induction that we can perceive causality. Causality is an entity acting in accordance with what it is, and we perceive this all the time because we don't perceive in isolated bits, snapshots, parts of things that we have to then work to see as wholes. We perceive entities over time, and what this means is that we perceive entities moving, doing things. When we see an animal walk, or a ball roll, we're seeing *something do something*. And we perceive that as a whole, it's not two things we perceive, the entity on the one hand and the action on the other, rather it's the one acting entity. Thus there's no problem of how we can come to see the relationship between what the thing is and what it does. Walking is moving a part of you, your legs, but, you know, you see a man and you see he's got these parts and you see he's doing this, and that's what it is to see him walk. And it's not big coincidence—how come it's the things that have those two appendages hanging down that move that way, it's a big mystery—it's part of what it is in seeing someone walk that you see that walking is a motion of the legs and that only a legged entity can do it, and so forth. Likewise with seeing a ball roll: you can see how the shape is involved in the rolling. You can't see everything about it, but you can see something about it. To deny that we perceive causality in this way would be to revert back to the sensation view, to disintegrate perception into a bunch of atomic sensations coming from different sense organs—to, for example, take apart the "heft" from the "hefty thing" you perceive when you pick it up. More generally (now putting together this causality point with the more general point about the problem with the sensation view of percepts) to view percepts as sensations or as groups of sensations

is to strip them of some of their content—is to strip them of the interconnected integration that they have; and this is going to make them an inadequate basis for knowledge.

1.3.3. Perception as distinct from perceptual judgement

The next thing to distinguish perception from is perceptual judgment. A perceptual judgment is a conceptual identification of the content of a perception. For example, "That's a man", "That's red"; these are perceptual judgments. People who don't confuse perception with sensation tend to confuse it with perceptual judgment. In general, people seem to think that there's just sensation and perceptual judgment and nothing in between them, and when they're talking about perception, they identify perception with either one of those two things: with an atomic sensation of red, or with the sentence, "That thing is red." But we need to distinguish these two things. Perceptual judgments necessarily, always, in every case contain content that is not in the perception, and they always omit some of the content that is in the perception. For example, suppose I were to hold up a blue cellphone and ask you, "What is this?" You might say, "That's a cellphone." Well, clearly this judgment, contains a whole lot that's not in the perception itself. For example it contains the knowledge of the relationship between this thing and old analog phones (which don't look anything like it). It presupposes knowledge of telecommunications, of language, of phone networks, maybe even of satellites, of cells by which the cellphones communicate. There's a lot of knowledge contained in the proposition, "That's a cellphone" that's not contained in the visual experience of seeing of a cellphone. And the same applies even if you utter a much simpler proposition about the cellphone when I ask you what is. For example, if you say "It's a blue thing." This proposition relates the cellphone to other blue things, to other shades of blue that the cellphone itself does't have, and it differentiates the phone from red things (even ones that aren't present in one's field of vision). None of this is part of the seeing of the blue cellphone. So as soon as you identify in conceptual terms what you see, you add content that's not present in the perception.

But you also take content away. They say a picture's worth a thousand words, and they're understating by a whole lot. Let's think of all the things that would be in the perception of a metallic blue cellphone that would not be in the propositions "That's blue" or "That's a cellphone". What about the exact shade of blue, for which you probably don't have a word? What about the precise way that the light reflects off of it, the exact contour of shading that you can get from it? What about the exact position of the cellphone relative to everything else in your visual field?

Perceptual judgment necessarily contains in one way more and in one way less content than any given perception. As we'll see later, all this extra content that's brought in comes ultimately from other perceptions—it doesn't come from nowhere, it's not something our mind contributes to the scene—but nonetheless, it's not part of the perception of the object that you're making a judgment about. Now, if you confuse perception and perceptual judgment, you end up thinking that the senses can deceive you, that they can tell you things that aren't true. For example, suppose the thing I held up and you thought was a blue cellphone was in fact a candy dispenser shaped like a cellphone. If your senses told you that it was a blue cellphone, your senses would be lying.

Now, of course, nobody thinks that our senses tell us things are *cellphones*, but they do think they tell us things like that it's *blue*. But suppose that the light was wierd so that when you got closer you saw that it was green. "Well, then your senses might be decieving you", would think somebody who held this view. Or, there's the famous "bent stick in water" illusion which Dr. Binswanger recounted for us the other day in his lecture. Isn't that a case when your senses are decieving you? They're telling you that the stick is bent, when it isn't. Well, your senses don't *tell* you anything, they don't engage in judgements. To get a good picture of what the senses don't do, think of the show *24* (which everyone should watch). There's Jack Bauer, the government agent who's always going into

dangerous situations, and he's got an earphone in his ear, and he's communicating with a woman called Chloe who's feeding him all kinds of information like, "Jack, there are ten hostiles around the corner." Your senses don't do that. Your senses aren't Chloe whispering to you about what's out in the world. Your senses just reveal the world to you, they just show you what's out there. To see is just to see, and you can't say any more about it because it's axiomatic; you just have to make the gesture of everything coming in. Perception is perception, it is not sensation or perceptual judgment. And it does not make sense to say that the senses decieve you.

There are other related arguments for the invalidity of the senses, famously the argument from the relativity of the senses, and there's a similar error here. I'm not going to go into it. Again, this is not an exhaustive treatment of each of these issues, it's an outline. If people want to ask about other arguments against the senses and what's wrong with them—that is, against the validity of the senses, other arguments that the senses can err or decieve us or whatever—you can do that in the question period, but for right now I'm not going to go into it. Rather, I'm going to now distinguish perception from a third thing: what I'm calling visualization.

1.3.4. Perception as distinct from visualization (which unlike perception is fallible)

Now this is not a point that's addressed by Ayn Rand anywhere. It's a point of mine, and, in a way, a point of Aristotle's. But I think it's important to understand more clearly what perception is and isn't, and in particular why perception can't be wrong, because it's important that perception can't be wrong—that it's infallible. This is why it can be the base of our knowledge, why we're secure in knowledge that we've built on perception. I want to distinguish, then, perception from visualization.

Visualization is the perceptual projection of how an object will act or respond, which is based on one's prior perceptual knowledge of it or like objects. Here's an example: an animal that jumps around in trees, say a chipmunk, has landed on branches of various sizes before throughout his life, and in each case he percieved how they supported him. This is again an example, by the way, of perceiving causality. Thing animal jumps and he lands on the branch and the branch bows a little bit and he feels it bow and he feels the rush of the wind and he feels it spring back; what he's doing is perceiving the branch supporting him. That is, he has a perception the content of which we would describe with the proposition "The branch supports him"; as we've just seen, the proposition imports more knowledge than is in the perception itself and omits some of the knowledge from the perception. Nonetheless, we can say the animal perceptually knows that this branch supported him. Now he's done this with a lot of branches before over the course of his life, and now here he is, about to jump onto a new branch for the very first time. Chipmunks do this; they jump onto new branches. How do they know which branches are going to hold them up and which ones aren't? Presumably they must do something like picture what will happen when they land on the branch. Now maybe they don't quite picture it the sense of we do when we picture things, but somehow they use their knowledge of what happened when they landed on past branches to form some kind of expectation about what will happen when they lands on a new branch, and so our chimponk is able to predict and expect, something about what will happen when it lands on the branch. It jumps, it lands on the branch, and all is well.

Except that sometimes all isn't well. Sometimes the branch snaps, the chipmunk falls, and the chipmunk dies. The chipmunk can be in error, or fail—there's something that can go wrong with what the chipmunk is doing in this case (assuming, that is, that chipmunk's can even do this, but I think it's very likely that they can), and we have to distinguish this from what's happening in perception, so that we can see how the kind of error that's possible in this chipmunk projecting case—in this case of visualization as I'm calling it—is not a problem with perception. We have to see how this is not a case of perception decieving the chipmunk, because if perception lies to chipmunks, it might lie to us, and it won't be a secure basis for our knowledge. Notice that the sort of failure or error that take places in this

case of visualization is possible even if visualization is deterministic, as I think visualization is in animals. So it's not an issue of free will or not.

The reason that the percept: the perceiving of the branch can't be wrong, but the projecting of what will happen when the chimpmonk lands on the branch can be wrong, is that the percept is the first entry into the animal's conciousness of the relevant material. The branch's shape, where it is, etc., everything it percieves about the branch just impresses itself on the animal's consciousness from without via the actions of a complex physiological system. The active process involved in perception is performed by the brain, and perception serves as the basic form in which the animal registers the way the world is. Visualization is a further processing of perceptual material which takes place within consciousness which aims at predicting the future—the near future. It aims at predicting that which does not yet exist, and which therefore cannot yet be perceived. And the visualization can be good or bad, right or wrong—I'm not really sure what words to use, but it can be judged this way by us—in accordance with whether or not the future turns out as the animal predicted. And the animal can discover that it's right or wrong (if he lives to see next jumping) by his future perceptions. So even if he's dying, the animal, as he's falling from the branch, has discovered:"Oh no! This branch didn't support me." (Again, this would not be in words, but in the form that animals can be aware of things.) So I think the reason why perceptual projection or visualization can fail, in a way that perception cannot, is that it's a processing of material that takes place within consciousness, whereas perception is the initial registering of the way the world is, against which other states can be judged.

Now some people, including some very good people who work on perception, call what I'm calling "visualization", "the perception of affordances". They think that what happens in the chipmunk kind of case is that the chipmunk sees—*percieves*—that the branch will hold it. But I think that this is a serious error for two reasons. First, this visualization process takes place within consciousness; it's not a process that merely gives rise to consciousness like the neural firings which give rise to perception. Moreover it depends on prior consciouss content in a way that perception does not because perception is the initial way in which we, at least as human beings, are aware of things, and: visualization can be wrong in a way that perception cannot. Now, you might think that this is cheating: I'm just defining things that can be wrong out of perception, but whether something can be wrong or not is an important difference. And there are other differences between this and perception, and that's a basis for making a differentiation. And it's a necessrary differentiation because we need to distinguish those states of our consciousness which cannot be wrong, which are infallible, from those which are fallible, because we're going to need, when we're at the conceptual level, to make sure that we're reducing our conceptual conclusions back to the infallible stage: to perception itself. We're going to want to make sure that our knowledge is built on that which is a direct contact with reality and which cannot be wrong. Perception is the automatic and basic form in which certain animals experience reality because their brains integrate sensations in a certain way. And we can see why its basicness means that it cannot be wrong.

Now there might be other such functions like visualization (or what I'm calling visualization): certain acts of memory, certain acts of association that maybe animals can perform, and it's not really important for our purposes what all these acts are. What is important is that these are cases of "post-perceptual" processing. They're rudimentary automatic versions of post-perceptual processing, and post-perceptual processing is a wider category of which conceptual cognition is the epitome. Conceptual congnition is going to differ from all these other forms of post-perceptual processing in that it is volitional. And the fact that it's volitional is going to give rise to all sorts of other problems that can happen with it, all sorts of other things which are impossible for processes like visualization. And it's these other kinds of failures that we normally mean by the concept "error". So I'm not certain what word is best to use to capture the way in which the chipmunk can go wrong. I don't think it matters that much whether we call it a certain kind of error or a certain kind of failing, but it is important to

differentiate it from the kind of errors that James Taggart makes, for example: that is, by being irrational.

1.3.5 Perception as our basic form of awareness

So we've distinguished perception from three things. We've distinguished it from sensation, we've distinguished it from perceptual judgment, and we've now distinguished it from visualization. I want to make one last point about perception. Namely (and it's a point I've really been making before) that perception is the base of all our knowledge. Now this point—that perception is the base of all our knowledge—is really not a point about perception. It's a point about the other types of knowledge that we have. Namely, it's really a point about concepts and propositions: that they're based on perception. And to see how they're based on perception, we have to study the concepts and propositions, which we're going to do in subsequent days. Once we see how they work, and we see that what they essentially are involves processing perceptual data in a certain way, we'll then see why perception is our basic form of awareness. But it's worth saying a little bit about it now in response to certain arguments that have been given against the view that perception is our basic form of awareness, because I think it does make a little bit clearer what perception is. Another way of saying what I mean in saying that perception is the basic form of awareness is that all awareness that is *not* perception is built upon perception. We have no independent kind of awareness—independent of the data that we get in perception. And the famous phrase that's used in the history of philosophy or in philospohy generally to describe this view is the phrase "*tabula rasa*", which means "blank slate". The idea is that man's mind is a blank slate which is written on by perception. The alternative idea is called the theory of innate ideas: that we have some ideas in us independent of perception. It's sometimes called "nativism" or "rationalism". (There are subtle differences between these two; we'll discuss rationalism later.)

Now, nativism is having a huge resurgence today. Psychologists are all excited that they think they've proven various innate ideas that we have. And it's worth looking at one of the things that is sometimes used as evidence for innate ideas (for built-in concepts or whatever you want to call them), because I think helps us to understand perception in a little bit more depth. It turns out that infants focus selectively on face-shaped things. If you put something that's shaped more or less like a face in front of an infant, it will look at it longer than it will look at other things. It will show some kind of an interest in it. It's not just a face but anything that's very crudely face-shaped—I don't know how crudely, but, you know, a bunch of dots in the right configuration, it will look at it.

Now it's sometimes supposed, on the basis of this, that infants have a concept, or a built-in category, or something like that, of faces. But this just doesn't follow, and the idea that it does follow is a remnant of a sensation view of perception. Why can't it be that our perceptual systems are simply better at discriminating face-shapes than they are at discriminating other shapes? We're, for example, better at discriminating colors within certain ranges than we are at discriminating colors within other ranges, we can hear certain sounds but not hear other kinds of vibrations, and we can probably discriminate sounds better within certain ranges of wavelength.

Someone objects: "But if we discriminate faces better than other things, then we're not really getting at reality as it is. Where faces are no more set off from the rest of reality than anything else, in reality, somehow for us they are. Wouldn't that be inaccurate? Aren't we putting them under concepts? Aren't we adding something to our perception?" No, that's just a form-object confusion. In reality, things aren't set off or not set off, things just are. And our consciousness discriminates them: that's what consciousness is, it's an act of differentiating and of integrating. And if we can differentiate certain things from the background better than we can others, well, so what? That's just part of perception. And this is true even if it turns out to be very important to understanding human

development. Maybe we couldn't, for example, learn language unless we were able to perceive faces in a certain kind of detail that we are only able to do because we're better at discriminating them. Again I'm just speculating that we're better at discriminating them. Even if it's that important and it turns out to be a major fact about human beings that one would have to know to have a psychological theory of human development of a certain complexity, it would remain just a fact about perception. Now I'm not a psychologist, I don't really know this research well, and I don't know that babies are better at discriminating faces. I don't even know that the research that babies look longer at faces is valid—I haven't read all of it. The point I'm making is just a point about the nature of these sorts of arguments.

Q&A

Christina: So you're saying that you can see cause and effect, I mean it goes...[incomprehensible]

GS: Right, yes

Christina: And, um, for some reason, I [incomprehensible] think that there is something more to it than just perceiving, ...you see [...] I can see you walking [....] fire burning

GS: Okay. So, the question is to clarify what I mean in saying that you can percieve causality—perceive cause and effect, and she gives some examples. She thinks she understands it in the case of perceiving me walking, and maybe in the case of perceiving pushing a ball, but what about perceiving fire burning a paper, for example? And along with this: do we perceive cause and effect like we perceive qualities of things? Well, first off I want to clarify: we don't percieve... it's not that the only way we ever know of any causal connection is by perceiving it, there are lots of them that you can't percieve. For example, you can't perceive, um, Kant causing Naziism... right? I mean you need... that's, that's, ...there are a lot of causal connections that you can only establish through complicated conceptual work. Um, it's that you can perceive certain ones. Um, now... and I don't even think you can perceive, always, everything about the causal connection, or that you can ever perceive everything about it—there's a lot more to learn about balls rolling than you can tell just by looking at them. For example, the relationship between the amount of impetus in your push and the way that they roll, and the direction that they roll, and how does their mass factor into it, and all of that is figured out by physics. Nor is the point just that you perceive the causality like a quality along its other qualities, like: "It's blue and it rolls." What you perceive is the thing acting. I think that part--that you perceive the action--isn't problematic. So, the point that you perceive causality is just that you don't perceive the thing and its action as two disconnected things, and then you have to wonder how they relate to each other. Rather, you perceive a whole, which is a thing acting over time. You percieve its action as part of an integrated whole, with everything else that you perceive about it, whatever that is. So now think about the case of perceiving the ball rolling: you perceive the relation of its shape to its motion: rolling is, after all, following its contour along a surface. Now, what about the fire burning? Well, what is burning something? For something to burn is for it to catch on fire, and if you put a piece of paper into the fire, you see the flames moving up onto the thing and consuming the thing. Um, you don't... So you can see something happening, you see a whole, um, that's, uh, involves the fire acting on the paper. That's what you see. Now there are two ways really, and they amount to the same thing: to look at what's involved in denying that you perceive causality in this way. It's really a disintegrating of the content of perception, and you can think of it as reducing perception to sensation, or you can think of it as reducing it to perceptual judgement in a certain way. Think of the "Chloe talking in your ear" model, you know, so Chloe you... the fire comes in front of your eyes, and the paper, you know, gets put into the fire, and Chloe says into your ear, "There's fire. There's paper. The paper's moving near the fire. Now the paper's burning." This is kind of Hume's view of things - he thinks of it as sensations, but, now... and then Chl... you know, you take another piece of paper and you bring it by the fire and Chloe

says, “Paper. Coming near the fire. It --it’s burning again.” You know, and then he goes, “Wow, there’s a constant conjunction,” you know. “Every time I take paper, then the next thing Chloe tells me,” you know, Chloe being your... “it starts burning.” But that’s not... it’s not like you get this in little atoms: paper, moving near fire, burning. You see a whole and you see... just think of what it is to see the paper start burning. You see the fire moving up it and start... you get all of that and that’s what it is to perceive causality, I think.

Um, I forget your first name.

Q: I was going to ask you, is that the exception... [ ]causality...

GS: Well,

Q: ...see, see the fire, and realize that something’s happening and [ ]

GS: Um, yeah, so the, the ... I...

Q: ...something caused it, but you may use your perceptual ...your [ ]

GS: That is, you... the idea that, the idea that you’re saying is that there’s a visualization involved in, in what I’m calling perceiving the causal connection. Um, I think that can be true in some cases, but then it’s not *perceiving* the causal connection. There’s, I think, ... I’ve described perception and visualization as distinct things, and they are distinct things, but I think they are, you know, sometimes continuous in our experience, it’s not like, uh, you stop one and do the other, you’re constantly having expectations based on what you percieve, and so forth, and it... I think, sometimes, there are cases where it’s not obvious which is which. It is, I think, if you isolate it and focus on it. I think you can always separate what’s part of the perception itself and what’s what I’m calling a post-perceptual process. Uh, so one, uh, uh, I, I want to make that caveat. Now, two, I think the difference between visualization and perception in this case would be--that is, the case of watching the paper catch on fire--is: as soon as you’re projecting what’s gonna happen, for example, that the next piece of paper will burn, or that the fire will continue up the piece of paper in a certain pattern, as soon as you do that, you’re visualizing, uh, or you’re else... you’re, you’re doing something conceptual if you’re a human being. What you perceive is not “fire burns paper.” What you perceive is “this fire burning this paper.” That is, perception’s always of particulars. Um, but what you perceive is *not* just a show of lights and then the show of light is on the paper. You know, if you think of fire in a... in the way a, a... sensations would, a kind of flickering. Nor is it the sentence, “Fire burning paper near fire paper burning.” What you perceive is the fire consuming the paper. Um, now even that, you can’t say what you perceive without *saying* it, and to *say* it is to put in the form of a judgement, and a judgement’s now something other than a, than a... a um... other than what you perceive, you know. But we... but, you know, think of... think of perceiving fire consuming a paper and,... and that’s what you percieve. Uh, and there’s a... you perceiving the fire acting on the paper. Uh, that’s how I, uh, understand it anyway.

Um, Harry [Binswanger]?

HB [faint on the tape]: Let’s not think too impoverished, also [...] it isn’t just yellow moves from here to there, you *feel* fire. You feel the heat, as you hold the paper, you feel your hand getting warmer and warmer, maybe burning.... probably drop it, as you see the paper first turning brown, and then burst into flame, what you feel is more heat, I mean there’s a whole... and you smell new gasses emitted as the paper catches on fire, there’s a whole [ ] of rich perception there relevant.

GS: Yeah, that’s, that’s very helpful. Um, uh, Dr. Binswanger’s pointing out that you recieve a lot more in perception than just kind of the color of the fire and then the paper kind of looking like fire. Perception is across modalities. You feel the the heat of the fire, you feel the, the, it... you know, the motion of the fire’s breath on the paper as it moves in your hand, you see, you know, all the subtleties of the progression of the fire up the paper, you smell the gas -- I’m repeating this because I don’t know

if it's getting captured. Um, so that's important.

Ray [Girn]?
Ray: [incomprehensible]
GS: mm-hmm
Ray: [still incomprehensible]
GS: Okay, so the question is: why do I think that the chipmunk, when it's about to jump on the branch, is visually projecting what'll happen when it lands, as opposed to just, in effect, feeling a desire to jump, uh, in a certain way. Well, for the sake of epistemology, it doesn't matter. That is, if what the chipmunk... if what happens to chipmunks is: when they get certain visual stimulations, they start wanting to jump, and then they jump, um, without having any what we would normally call an expectation of what'll happen, then there's no argument that comes up from, uh... to the effect that maybe perception is getting it wrong here. So there's nothing... there's no phenomenon in that case that we need to distinguish from perception. However, I think it's plausible that what happens is something more like what I described, and if that's true, that is, if what happens is something more like I described, it's important to distinguish it from perception. So, it's, um, for our current purposes, that visualization exists is not important, and that's no part of epistemology. What is part of epistemology is, in effect, that *if* visualization exists, it's something distinct from perception. Now why do I think it exists? It's a rather complex information coordinating, that it seems that the chipmunk has to do in order to do this. Um, it's not clear to me that it could happen in this way that you're describing it, that is, just by it has a certain stimulus and then it has a certain feeling, um, it has to at least be able to use the information from past perceptions to generate that feeling, uh, and it seems very plausible that it would do that in the way that I'm suggesting. There's an analogous function in human beings, uh, of visualizing the future, and it doesn't always take place volitionally even in human beings, say if you're a phobic, um, and you're kind of near an edge and you're afraid of heights, you might be able to... it might happen that without being able to stop yourself, you just start, you know, seeing youself plummeting to your death, or something. Um, so I think there's, um, the... the complexity of what they're doing and the analogy to the kinds of things we do in those kinds of cases. Um, supports at least a ...uh, uh, this hypothesis that this is what's taking place, and I consider it in my mind at the level of a probable hypothesis, because I'm not a uh, a scholar of animals. Um, but, I bring it up here not to argue for its existence, but to say that if there is such a thing, it oughtn't be confused with perception. Because I think that it can be confused with conception... perception easily, and often is.

Um... Allan [Gotthelf]?
AG: ...infallible, and I'd like to ask about that. There's a phrase you mentioned [...] was basic, while visualization was not
GS: Right
AG: Second point was that perceptual [...] was physiological but visualization is a conscious process
GS: mm-hmm
AG: so, my question is, uh, how are those two issues related, and ... more fundamental than the other, but I'd like to ask a question just in this context: that you said visualization can go wrong.
GS: mm-hmm
AG: Now, in what way can it go wrong? Well, I... my notes, it's perceptual projection of how an object will act or respond. Well, there's a *goal* that seems to be involved - an action-related goal - if I'm a squirrel or a chipmunk and I project the branch as there and it's a tiny bit away from there [...] it's irrelevant, right?, so if it holds me versus not, that's tied to my jumping. Now, so, is the goal-oriented nature of visualization tied to its ability to make a mistake? And if so, what about the goal or--

the goal oriented feature of perception?

GS: Okay, so there are sort of --see, if I have to repeat this all, I have to see if I can hold it, um... okay, um, so There are two features that I stressed in talking about visualization as opposed to perception and why visualization is fallible. First that visualization is post-perceptual, that is that it's a process that takes place and its input is perception, and second that it's in consciousness, whereas the processing that gives rise to perception is physiological. Um, so what's the relationship between those two points, and how do they relate to one another in showing that the one can go wrong, that is, um, uh, visualization, wherease the other, perception, can't? Um, the other point was... well, let me answer that one and then... oh, oh, yes, the fact that visualization is action-directed: is that why it can go wrong: because it projects what'll happen when you take a certain action, whereas perception doesn't, and yet perception itself is for the guidance of action, so if it was the action-directed character that made it fallible, then that would make perception fallible, too, perhaps. So these are the two points. First, I want to give a qualification: I did use the phrase "go wrong", and I don't really like that phase, and I don't think I should have used it. It was something that occurred to me as I was speaking, it's not how I said it in my notes. What perception and what visualization can't do if it's determini ... if it's a ter...deterministic process is precicely "go wrong." That is, be conducted in the wrong way. Because the way that its conducted can't be judged: it happens automatically. What can happen, though, is it can cause the animal to expect something, and that thing that it expects to not happen. So we can have something that's analogous to falsehood or error. Exactly how it relates to falsehood or error is a difficult questions; I have some thoughts on it, but I'm not going to say that just now. Now the issue of how does the post-perceptual um, feature of it, of, uh, visualization relate to the fact that it's done in consciousness and not done physiologicallly as I'm saying that perception is. Well, I think they're two ways of saying the same thing. That is, you get content into consciousness by some physiological processs - the one that gives rise to us seeing - uh, then you process that content, and you're processing it qua conscious content, and that's necessarily *not* physiological because it's something that takes place in consciousness, though, of course, anything that takes place in consciousness, or, if not "of course," strongly probably anything that takes place in consciousness has some kind of physiological component or correlate or thing underriding it. But I think it's essential to what visualization is that it's something that's done with conscious content. Whereas it's essential to what perception is that it is something that results from non-conscious processing that gives rise to the, so to speak, first or basic conscious content. Now, I don't think that what's essential to this difference that gives rise to error or not, or the possibility of failure of this kind or not, is the action-guiding factor here. All conscious processes are teleological in the sense that they aim at identifying things. So in that sense, they're all active. And I don't think it's in particular that visualization... the visualization is about what'll happen when the chipmunk jumps on the branch. I mean, you could suppose that there are two chipmunks and the first one jumps to the branch, and the second one is sitting on the previous branch expecting the branch to hold the first chipmunk, and then the branch breaks and this second, spectator chipmunk is as suprised as, you know, as anyone that this happens. I think that, you know, he wasn't contemplating jumping, so I don't think that the chipmunk's, you know, conteplating action is what's essential.

Now, Betsy's been trying to get a word in edgewise.

BS: Okay, okay, um, two points about what you call "visualization." Can you reduce that to two automatic processes, namely the association of past actions with their results, and the retaining of that memory.

GS: That is, Betsy asks, "Can I reduce vis... can visualization be reduced to (what I'm calling visualization) to the association of past actions and their effects, and the retention of the memory of the past actions

BS: Yeah

GS: Um

BS: And the results, so that when you see this, now, a branch of a certain kind, you associate that with successful landing on the branch.

GS: Um... surely... I mean, obviously memory is essential to this process, right, remember past branches. Obviously it's essential that you saw um... you perceived the results of the previous actions. Um, I think... And now that you're able to associate the new branch with the previous branches you perceived, and thereby new actions that you might take with this branch to the old actions that the old branch, um, the way the old branch responded. Um, I think there's one more step to it than that in this case, that is, the expec... the forming the expectation, but if not, not. I mean, I don't think it matters all that much.

BS [beneath loud water-drinking noises]: Okay, great, could it also be that there isn't really any expectation and that it's something that you're reading in, that in fact all the chipmunk is saying when he jumps is "I did it in the past." And that's true, and that's infallible.

GS: Uh, no, I don't think... I don't think that's possible, but if it is possible, it would be sort of like Ray's question over again, that is: if there isn't such a process that's so distinct from the other processes, um, then there's no need to distinguish it from perception. So again, my distinguishing it from perception is not supposed to be some exciting new thing, it's supposed to guard against a possible error. Now, I ... that doesn't seem right to me, that all you're doing is thinking about the past cases, because there has to be a step from the animal's associating about past cases to initiating action in the new case.

BS: But, could it just be that it...

GS: Well, let's, let's stop on this now, um, there's time for, I think, one more question. Does anyone have that's not on this issue, because I don't want this issue to monopolize it if other people..

...okay, Tom?

TB: I wonder if you could clarify the correct and incorrect views of how sensations give rise to perceptions, because you, you said at one point, perception *is* a group of sensations that are automatically integrated [rest incomprehensible]

GS: Okay, I've said that uh, uh, uh we can't reduce perception to sensations, that is, and that... part of that is we can't think of a perception as *just* a group of sensations, but I also said that a perception is a group of integrated sensations or an integration of per... of sensations, and um, Tom want's me to differentiate those two, what's the difference. Um, uh perception is a form of awareness that is an awareness of three-dimensional entities--of single, whole things persisting over time--and that form of awarenss comes to be by the integration of what we might call sensations. I'm happy calling them sensations, but they're not sensations in consciousness, they're not atomic awarenesses of things, but they're rather, basically, currents taking place in our nerves. I mean, they're the, you know, nerve firings or whatever they are, they're some non-consciouss thing. So that's um the... the perception is a unitary form of awareness of an entity, and it comes to be by some process of integrating signals coming into our sense organs. The wrong view is that a perception is in effect a set of sensations, it's, you know, the silver of this pitcher, and the coldness of this pitcher, and the shape of this pitcher, and the shininess of this pitcher, and the jingly noise of this pitcher all in a "bundle" to use a word that I think Bishop Berkeley used (and if he didn't use it, it's become associated with him). Percepts are not bundles, they're unitary forms of awareness of unitary things, of entities.

We're out of time, and I was told that I'd have to be very careful about finishing on time.

Lecture 2: Concepts

2.1. The Problem of Concepts

Yesterday we were talking about perception. We were talking about, first of all, epistemology and consciousness in general and then about perception: how it can serve as the base of knowledge. Because of some questions that I got after the class, I want to emphasize here that what I'm doing is an outline, and that was especially so in the case of perception: I was not trying to answer *every* objection one can make to perception or every mistake one could make about it, but certain ones that I think are especially important. People who are interested in other cases can ask me about it in the question period or after class. But now we're moving on to concepts and to conceptual knowledge.

2.1.1. The role of concepts in knowledge

Our knowledge is held in the form of propositions, and the propositions are made up of concepts. This is why concepts is the essential issue. In simple cases, at least, propositions have a subject-predicate form. A proposition is the form: S is P. Cows are animals. Black is a color. Dogs bark. Subject predicate. Parrots are birds. Birds have beaks. There are more complex propositions which we won't be talking about (people can ask in the questions if they want), but this is the basic form, and it's the form that's most important to understand. Now, the knowledge that we hold in forms of propositions with the subject-predicate form is not limited to simple things like "Parrots are birds." We know that communism is statism, that statism is a pervasive initiation of force, that force paralyzes the mind, that the mind is man's means of survival, that the requirements of man's survival are the standard of value, and therefore that communism is evil. This is important knowledge, all in this subject-predicate form, all in the form of concepts. But what is a concept? I'm going to give a generic definition of it--that is, a definition that's formulated so as to allow for the possibility of different, competing theories of concepts, so that we can see what it is that the theories are competing about, what they're alternate explanations of.

2.1.2. Generic definition of "concept"

A concept is a unitary cognition (of the sort expressed by a word) of or applying to indefinitely many differing objects. For example, the concept "bird" refers to all birds, and birds differ from one another. When one thinks, for example, "birds have beaks", one thinks this *at once* about all birds. So the concept "bird", expressed by the word "bird", is a form in which we think at once about all birds, and it's a unitary form in which we think about all birds. Let me clarify what I mean by "unitary form" by constrasting it with something. I'm aware at once, say of everybody in this room—my visual field includes all of you; however, I do not have a *unitary* awareness of all of you, because one part of my visual field is Betsy, and another part is Eric, and so forth. Now consider the concept "bird"; it does not have parts, one of which is "roben" and the other "hen". It's not that the "b" in bird stands for robin and the "i" for hen, and so forth—those aren't even the same first letters. [laughter]

In Objectivism, we call each of the things which fall under a given concept—say, each bird that falls under the concept "bird", a *unit* of the concept, and I'm going to use that terminology as I proceed. I'll explain where it comes from, why Ayn Rand uses it that way later. But for now, a unit is one of the indefinitely many things of which a concept is a unitary cognition.

2.1.3. The nature of the problem

Concepts, and therefore propositions, are universal. That is, they're one cognition of many objects. Universal is just a Latin word you get by smushing together *unum versus alia*: "one over many". It was coined as a translation into Latin of a word Aristotle coined in Greek by smushing together the words that mean "of a whole" (for people who want to know the etymology: *katholou*). The concepts are universal, so conceptual cognition is universal cognition: unitary, singular; it is one cognition of many differing objects. And this brings us to the problem of concepts, or what's traditionally called the "problem of universals". In essence, the problem is the question: How do we have unitary cognition of what are, to perception at least, indefinitely many, differing things? How can we have unitary cognition of indefinitely many differing things? It's the questions of how conceptual cognition works. This question is essential to epistemology because conceptual cognition isn't automatic. One does not automatically identify communism as a unit of statism, or as being evil, even if one might automatically identify a given robin as a bird—at least if one's learned that concept. We have to know how we *can* do this, how we *can* think unitarily of indefinitely many differing things so that we can know how we *should* do it. How should we classify Pol Pot's dictatorship, for example, or President Bush?

2.1.4. Traditional solutions

Now, the traditional solutions to the problem of universals, and indeed the traditional way of thinking of this problem—the very calling of it "the problem of universals"—is an instance of the form-object conflation that we discussed yesterday. People do not normally ask, as I think they should, "How is it that we know objects universally?" Rather, they ask, "What are universals?"—that is, "What are the universal objects of which we seem to be aware?" They're doing the same thing as I would be doing if, after seeing a man blurrily, I asked "Who is this blurry-man that I see? Is he out there in the world, a blurry guy, is it just in my head so that I'm not aware of any man, but only of blurry thing in my head?" The traditional debate about universals follows the same pattern.

Universality is a form of awareness; traditional philosophers tryed to make it into an object. Or, rather, they just took it as an object in their posing of their questions about it. We can see the results of this in their answers to the question. One school, the "realists", affirm the existence of "real universals" and they say that it's these universal objects that we actually know when we have concepts: a concept is an awareness of a universal object. Our knowledge of this universal object applies in a derivative or defective way to the various units of the concept because of some relationship or other in which those units stand to the proper object: the real universal. There are two influential versions of this view (there are actually some other versions, but two historically significant versions): the first is Extreme (or Platonic) Realism, where the universal is supposed to be a supernatural sort of archetype, of which the particulars are poor imitations. The second is Moderate Realism, the view often ascribed to Aristotle (I think incorrectly), on which the universal is an internal essence, an element *in* each unit which is identical to a corresponding element in each other unit, such that if you think about this element alone, ignoring all the things that are co-present with it in each of the units, the thought is equally of any of the units. We'll talk a little bit more about that in a moment.

The schools opposed realism are Conceptualism and Nominalism. We don't need to focus on the differences between these two right now, but what's in common between them is that they deny that there are any universal objects, and think that either there's some kind of fiction going on in conceptual thought, becasue we do have a universal functioning in out thought but it does not correspond to anything in reality. Either, on the Conceptualist view, the universal is a concept that's just made up; it's a kind of fictitious existent. On the Nominalist view, universals are percepts that serving a special function in some way. If people want to know about these two views and how they relate to each other, that's again something you can ask in the question period, I'm not going to dwell on it because it's not

essential to Objectivism. What is important to notice is that both of these non-realist schools deny that conceptual cognition constitutes knowledge of a mind-independent reality. They either revert to explicit skepticism, which is what Nominalists usually do, or they claim that conceptual cognition is knowledge, but it's knowledge of objects created or structured by our own minds, which is the position Conceptualists usually end up with.

2.1.4.1. Traditional solutions' lack of methodological guidence

The bigger problem, I think, with these views, though, is that none of them offer us any guidance on how to form or use concepts, because neither of them thinks there's any *how* to forming or using concepts. They think there's a universal object, either internal or external, not a process of universally cognizing particular objects, which process we can then get guidance on how to perform.

Well, let's concretize this lack of guidance by thinking about the individual schools. Start with Extreme Realism. There's a Platonic universal, a supernatural archetype. And how do we know it? Well, we just *intuit* it, we just automatically know it. Now, Plato does give us some methodological advice, but it amounts to preparing yourself to receive this information. It's what I used to call the "carrot eating epistemology". Carrot eating is the only advice you can give about perception: if you eat a lot of carrots, your eyes get better and then you'll see better. Well, likewise, according to Plato, if you do certain things, train your mind in a certain way, you'll come to just intuit the forms. There's a little bit more to it than that—you can refer to my Plato class if you're interested—about how you can clear false views out of the way, but essentially there's no positive advice on how to form and use concepts. All you can do is test whether you formed concepts, and are applying them, correctly, and this amounts to testing whether or not your beliefs are coherent, are consistent, form a whole. And it'll be significant that this is the only kind of guidance he can give you. But there's no guidance on how to make them consistent or form a whole, or how to make them get at reality.

Conceptualism does the same thing. According to it, you just form your concepts arbitrarily. So again there's no advice except that they tell you try to make them coherent, or else elegant, or else to fit your moral values, or some other standard, which standards aren't, and aren't supposed to be, derived from the needs of getting at a mind-independent reality. So again, no guidance in how to cognize conceptually, no guidance on how to conceptually achieve awareness. They might tell you how to shuffle around your mental contents, but no guidance of how you *need* to do that if you want to get at awareness of the world. It's not that the guidance is false, there's just *no* guidance, period.

Nominalism says the same thing as Conceptualism, except it's a little more cowardly, it tends to council caution: don't use to many large abstractions, and so forth. Again, we can talk about them later.

A more general point about these three schools, and really all the schools: we all implicitly learn how to form and use concepts as children. And as a result of this learning we've done in childhood, and all the conceptualizing we've done on the basis of this learning, it just *seems* right to us to group certain things together and to identify them accordingly. It feels right, they seem to us to resemble each other or to be similar, though, of course, we don't know how to do this in a given case. And other times, we find that what seems right to us conflicts, so that we know there has to be some sort of an error. It's because of this, that is, this non-omniscience, and non-infallibility, that we need a theory, and these schools give us no advice. What they tell us is to rely on what seems right to us until either we draw a blank or come up against a contradiction, and then to tinker, in some unspecified way, with our mental contents, until we either feel better or have a supernatural insight. None of these schools tries to get behind our inchoate sense that certain things belong together, and to articulate the basis of this sense so that we can take conscious control of the process of conceptualizing, so that we can check the identifications that seem right to us and form new identifications in the cases where we draw a blank.

2.1.4.2. The false theory of "context omission"

The partial exception to this is Moderate Realism which *does* have a theory, it's just a very inadequate one. Remember, Moderate Realism is the view that a concept refers to an identical element present in each of the different units of a concept. So it follows from this theory that we can form or validate a concept by looking at a bunch of things and seeing if there's any identical element we can find in them. If there is, then we can have a concept for them referring to that element. And if you're not sure whether something new is a unit of that concept, you can just look and see if it has the identical element. So there *is* a kind of method here, or at least there promises to be one.

For example: if you look at three birds, the first one will be green, it'll have a beak, it has feathers, and it flies. The second is blue with a beak and feathers and it flies. The third is pink with a beak and feathers, and it doesn't fly. All they have in common is beak and feathers, so beak and feathers is the essence of "bird"—the thing that all the birds have in common. Birds differ from one another by the fact that in some birds the beak and feathers are accompanied by one color, in other birds by another color, in some birds they're accompanied by flying, in other birds by non-flying. We can then say that there's a theory of how you can form concepts here: the theory of picking out the common stuff and getting rid of the other stuff. And this theory, I think, has been very aptly termed, by a scholar of John Locke, the theory of "context-omission". And that's the term I'm going to use for it.

Why "context omission"? Because you focus on the things that are the same in all the different birds: the feather and the beak, and you omit the context: that is, the other traits with which the feathers and beak happen to be present—different traits in different birds. So, the Moderate Realist says: we form concepts by context-omission, by looking for the identical element in a number of things, and focusing on it, and ignoring the elements in which the things differ. This theory of context omission is, to my knowledge, the only substantive answer ever given to the question of how do we form concepts (that is, prior to Ayn Rand). And all the schools of concepts fall back on it, at least sometimes, even the Realist... even the Nominalists and the Extreme Realists all somehow, in different cases, fall back on it, in way or another, when pushed.

The theory is false. You just cannot find anything identical like that in birds. The things that we called the same features in the birds that we listed, that is, the beak and the feathers, are only the same in the same way that the birds are the same as one another as wholes: in the sense that we think of them as units of the same concept. Beaks are all identically units of the concepts "beak", but each beak differs from other beaks: some are harder, some are softer, some are longer, some are more or less curved, they're different colors, they're in different relative positions on the bird, and so forth. And likewise with the feathers. If you try to *iterate* the process of context-omission—that is, if you try to say, "okay, well, we'll just perform it on 'beak' now, to get at what's the same about all beaks", and so on, you'll eventually end up with nothing at all. You would end up focusing on a single attribute of birds, for example, maybe, the length of their feathers, and then you'll see that the length of the feathers even differs from bird to bird. So context-omission does not work as a theory of concept formation, and along with it, Moderate Realism does not work as a theory of concepts. Or, if one wants to retain Moderate Realism, one has to make it more like Platonic realism than it seemed at first; one has to say that we know essences by intuition, because they can't be the sort of things that we can have a method of discovering.

2.2. Measurement-omission and its presuppositions

Ayn Rand replaces context-omission with measurement-omission. She authored a new theory of concept formation. And to illustrate what measurement-omission is, and why it's different, it's helpful

to begin with the case that we left off with with the Moderate Realists. Namely, the case of the lengths of the feathers of different birds. Notice that Ayn Rand introduces it with the example of the concept "length", focusing on a single attribute. If one takes the different lengths and looks at them—looks at the lengths of a sparrow feather, a pigeon feather, a hummingbird feather, a parrot feather, an ostrich feather—one will not find anything left after one removes the differences between the lengths that one can call length. There's no universal "length" or universal "feather-length" that can be distinguished from the specific lengths of each of the feathers. You can't omit the context in which the feathers have their particular lengths.

2.2.1. The nature of difference

But notice that the lengths do stand in a relationship to one another. The lengths *differ*. To differ, things must be comparable. They must, to use a phrase from Aristotle, differ in the more and the less. Think of the comparative form of adjectives: longer, bluer, more yellow, etc. Things have to be comparable in order to differ. That is, and least the way I'm using the word "differ" now; you can use what word you like for it, but there is a relationship that's captured by "differ" used in this way, such that two lengths differ from one another but length and color don't differ from one another, because length and color can't be compared. They're not comparable; they're disparate from one another; to use Ayn Rand's word, they're not *commensurable*.

The identifications of relationships between things, of differences in the more and the less, are actually a form of measuring things against one another. If you say, "This feather is longer than that feather", you're measuring one feather in terms of the other. You're getting a sense of how long the one feather is by comparing it to the length of the other. And likewise when you say, "This is bluer than that", or "That's redder than that", you're measuring one against the other, and therefore the must be *co-mensurable*: measurable in a common unit against a common thing. For things to be different, they must be comparable or commensurable, and what this means is that there is some *respect* in which the things differ. Aristotle makes this point nicely, incidentally, when he says, "That which is different differs *from* something *in* something, so that it's necessary for there to be something the *same* in which they differ." That which is different differs *from* something *in something*, and that *in something* is the respect in which the things differ. It's the hue, in the case of the difference between red and blue, which differ in hue, or the length in the case of the difference in the length of a ostrich feather and a hummingbird feather, which differ in length.

We can think of this something, this respect in which the things differ, as an axis. When such axes are used in the process of forming concepts, which we'll turn to shortly, Ayn Rand calls them conceptual common denominators, or CCDs. And we'll see her definition of a CCD shortly. Now, why CCD? Why conceptual common denominator? Well because it's like a common denominator in mathematics. Try to compare 42/63 and 55/99. Which is bigger? Can you compare them to each other, can you measure them against one another? Not as such. But if I tell you that they're equal to 6/9 and 5/9, you can. Six ninths is more than 5/9. You can measure them against one another because they have a common denominator. Likewise you can measure blue and red against one another. You can see the blue is more towards this side of a line than the red is, only because they have a common denominator: hue. So, difference, the relationship in which the different lengths of different birds' different feathers stand to one another, is a specific relationship, and it requires commensurability, it requires that the two things vary in the more and the less, that they vary along some axis of measurement, along some conceptual common denominator. That's what it is to differ as opposed to merely be disparat. (This point about the contrast between difference and disparateness I get from Dr. Binswanger, although it's also in Aristotle in different terms, and in Ayn Rand.)

2.2.2. The nature of similarity

Now let's turn from the nature of difference to the nature of similarity, which similarity we'll see is part of the basis for a concept. Similarity is the relationship of comparatively slight difference. Similarity is a type of difference. It's a difference viewed relative to some other, larger difference. That requires that two things can only be similar in comparison to a third thing; we'll call that thing the foil. So two blues cannot appear similar to one another except when contrasted to a third, more different color. Now Carl Barny and his shirt aren't here this week—they were last week: he had this great shirt that had two almost identical shades of blue right next to each other, and then under them a big patch of yellow. And if you just covered up the yellow, you saw two different colors. But if you bring the yellow in, then you see two shades of blue—that is, two distinct shaded that looked alike as contrasted to the yellow. We can all imagine this; we don't need Carl's shirt, but it was really vivid last week. Okay, similarity is comparatively slight difference, which means it's a three term relationship, a relationship between two things which are similar, and a third thing which those two first things are similar by being more different from this third thing than they are different from one another. The similarity is the comparatively slight difference of the two things as opposed to the third, and all three of these things must be commensurable—that is, they must fall along a CCD. In the case of color, the two blues and the yellow fell on the CCD of hue.

Now we're ready to define CCD (conceptual common denominator), and here's Ayn Rand's definition: "The characteristic(s)," (she puts 's' in parentheses to tell you that there can be one or more), so the characteristic(s), "reducible to a unit of measurement, by means of which man differentiates two or more existents from other existents possessing it." To see the blues as similar, we jointly differentiate them from the yellow along the CCD hue. Now, notice she says it has to be reducible to a unit of measurement, she doesn't say it has to be reduced to a unit of measurement. That is, it's importance for this process that there is a unit in terms of which one can measure the different hues, but to see things as similar, you don't need to actually do that, you just *see* them as similar, at least in the case of the colors. But what that seeing them as similar *is* is a form of perceiving them as insignificantly different from one another compared to their difference from something else along some axis which is reducible to a unit of measurement, though you're not perceiving its reducibility to a unit of measurement.

2.2.3. Units and the "unit perspective"

Now we can come back to the concept "unit". "A unit," writes Ayn Rand, "is an existent regarded as a group of two or more similar members." Notice it's an existent *regarded* in a certain way. It's an object insofar as you're aware of it in a certain form. It's not a blurry-man, but a man seen blurrily. A unit is an existent regarded as a group of two or more similar members. It is an existent thought of as part of a group of existents, which are jointly differentiated from a foil along a CCD because their differences from one another along that CCD are dwarfed by their differences from the foil. Notice, then, that nothing is a unit in and of itself, nor is a unit anything that anybody feels that it is. Rather, it's an existent that's regarded a certain way because it stands in a certain relationship with other things, which relationship we know about. It's an existent regarded as a member of group of similars, where we can regard it that way because we can perceive its similarity.

To view something as a unit is to take a certain perspective on it: the unit perspective. It's to think of *this* shirt as like *that* one. *This* one as being a member of the group that comprises that one and that one and that one. You're thinking of it as "one of those", where "those" are like each other in some way. Now, there are certain advantages of doing this. When you encounter another blue thing, you can say, "this one is like that one and that one, in that way that that one and that one are unlike that one."

But that's burdensome. It takes a while; it's hard to hold it in mind. We have, now, a cognition of one thing in a complex relationship with a bunch of other things. What we don't have, though, is a unitary cognition of many differing things. And because of that—because of this disunity, this kind of complexity of the perspective we have on things when we're just viewing them as units in this primitive way—the awareness that we have is fleeting: you just can't hold it, or you can't hold it for very long.

2.2.4. Integration by measurement-omission

To get beyond this, there's a process, and this is the process by which concepts are formed, according to Ayn Rand: the process of measurement-omission. We form a unitary awareness of the group of similar existents by omitting their particular measurements. This means, in effect, projecting the range along which the units fall. Instead of a bunch of dots near one another, we can think of a line segment. So if you can imagine—we said there has to be an axis for there to be difference, think of the horizontal dimension of this white board as the axis. You see the two blues represented by two dots that fall along the axis. The difference between the blues is represented by the diffeent locations of the dots along the axis. Now we add in the yellow over here—also along the same axis, but further away than the blues are from one another. Now, the blues look similar as opposed to the yellow, and we can see how we can add in more blues in the vacinity of the blue dots, which would be part of the blue-group as distinguished from the yellow. Now we can say: "this one is like this one as opposed to that one, and this one is like this one and this one in the way that they're different from that one" and so forth. The measurement-omission consists in connecting the blue dots to form a line-segment. Notice that I've drawn the line segment in suich a manner thant the individual dots are drawn over—eased. We now have a single thing: the range. And notice that the measurements are omitted. What that means is that the specific place within the range is omitted. As Dr. Binswanger puts the point, and I found this very helpful, "measurement omission is measurement inclusion." It's not like we're omitting the context of one of these dots in some way, like the Moderate Realists would. Rather we're including all the other dots. But that inclusion of all the other dots is an omission of the specificity of the location of each dot, because in drawing this line, we've blurred out the differences between the dots. We now have one line-segment encompassing a range, rather than having separate, discreet awarenesses of many dots along that range. That's measurement-omission.

There are two significant consequences of it. First, we no longer have to hold each unit separately in mind. Second, we've now already included every possible unit, because the line-segment includes everything along it. Measurement omission is the basic act of concept formation. It integrates our awareness of many similar existents into a singular, unitary awareness of them.

2.2.5 The role of words and how concepts are "mental entities"

There's another step, though, that we need to talk about. It's the introduction of a word. We introduce a word to stand for the concept. And the perspective of each unit as falling somewhere within the range of omitted measurement becomes associated with the word so that, the word can serve as one unit which enables us to grasp the whole range. As a result, instead of saying, "this one is similar to that one and that one, in the way that that one and and that one were different from that one," we can say, "this is blue." The word becomes associated with the perspective that we previously had to articulate by that complex "like this one in the way that they're different from that one and and that one." The word condenses this and makes it manageable and, therefore, retainable. The result is that the concept which is associated with this word is an integration or a new mental entity—both phrases Ayn Rand uses. She defines a concept, as we'll see a little later, as an integration, and she describes it as a

new *mental entity*. Mental entity is an analogous, or extended, usage of the word "entity".

An entity is something three dimensional out in the world that moves around and exists over time. There are not mental entities. What, then, is the sense of this extended, or analogous, use of the term? Mental entity is not a species of entity, a subdivision of entity. So what is it? Well, a mental entity is a mental state or process that's entity-like in that it's unitary and enduring. I think those two features are the important ones. It's unitary and enduring. Compare mental entity to legal entity. Microsoft is a legal entity. ARI is a legal entity. ARI is unitary and enduring in a way that the various people who comprise it are not. Before there was ARI, all the people who are on the board of ARI, or most of them, were around. There was Leonard Peikoff and Mike Berliner, Peter Schwartz and Harry Binswanger, and so forth, and they were all working individually, and, you know, in alliances *ad hoc* with one another to advance Objectivism. And if one of you wanted to give them some money to help, you could have. But what happened when ARI was formed? Well, now there's one unitary, enduring thing. When they die it can live on. And besides that ARI has a unity that the *ad hoc* aliences lacked. You contribute to ARI, and ARI can make decisions and enter into arrangements as a corporate body. It's one thing. Likewise with Microsoft: both are legal entities. It's on that model—the making of a collection of things unitary and enduring, and therefore able to act as one (and we'll see how that part is important in a moment)—that we should think of mental entity. A mental entity is a unitary and enduring mental state that can relate as one to other mental states and so function in thought as a unity. It's an integration, a unification, of many different things that enables them to function as one, to relate to other things as one, holistically, rather than bit by bit.

Another way to think of this is that a concept is an institutionalization of a perspective. We talked about the unit perspective, which is a complex act of thinking of things in relation to other things. Well, to form a concept is to institutionalize that perspective, to make it unitary and enduring over time. Think of the Ayn Rand Institute. And institute is group of people organized so as to function as a whole over the course of time. The units of a concept function cognitively as a single whole or unit in our thought. Yet another way to make this point—a little bit different, not now stressing the unity, but stressing the ongoingness and stressing the active nature of concepts—to form a concept is to put oneself on a *policy* of viewing things in a certain relation. And, I think, this point, and its relationship to the entity point that's established by the word, is captured by the quote Ayn Rand uses to capture what a child's (wordless) thought process would be in forming the concept length: "I shall identify as 'length' that attribute of any existent possessing it which can be quantitatively related to a unit of length, without specifying the quantity." This is a statement of policy. And what the child is doing in forming the concept is putting himself on this policy. He's able to make it unitary and enduring through the word "length." "I shall recall and regard anything as length which functions in this way; I'm on a certain policy with regard to identifying and relating things through the use of this word 'length'." That's what a concept is.

2.2.6 Concepts as future-looking

Notice that inherent in this is that concepts are future-looking. A concept is like a policy or a commitment. It's like forming a file. Ayn Rand analogizes concepts to files famously, and the point that most people stress in talking about this analogy, is that the file organizes and condenses data. And that's certainly true, and it's certainly significant—I think was the main point that she had in mind. But there's another important implication of this that I think is worth stressing. A file, if you have a filing system, does not only organize and condense data that one already has, it does so on the premise of keeping up with this method of organization. If one stops keeping up with it, it ceases to be the file that it is.

Suppose that you establish a "receipts" file for your income taxes by putting a couple of receipts

that you have in your pocket into the file. That file is not really your receipts file unless you consistently file your receipts in it over time. If you don't do that, it stops being your receipts file and becomes simply an assortment of random receipts that happened to get deposited there. It can't be used as your receipts file, say, when you go to do your taxes because it doesn't have all your receipts in it. The file isn't just the object—the folder with a bunch of things in it—it's the object insofar as it functions a certain way because of a policy that you're mainaining. The concept is like this kind of policy.

Likewise, if you form the concept "blue" by integrating a couple of shades that you've seen in the past, and later refuse to identify any newly seen shades as blue, what you formed would cease to be the concept "blue". It would cease to be a concept altogether. It would be, you know, "a list of colors I once saw." Okay. So concepts are future-looking, and that's going to be important. To form a concept is to commit oneself to continually cognitively acting, functioning, in a certain way across time. That's really going to be essential to understanding some things about objectivity later.

We've been talking about concepts as the result of the process of measurement-omission. Measurement-omission is the process by which we form concepts. There is more to understanding concept formation, however, than measurement-omission. There are more requirements for a valid concept than merely the commensurability which measurement omission presupposes. Commensurability and measurement-omission are necessary conditions on the formation of a concept, but they're not yet sufficient. We'll see what the other conditions are shortly. That is, I don't even think you can perform measurement-omission just anytime there's a commensurability. We'll see why in the next section.

2.2.7. Measurement-omission contrasted with context-omission

Before we turn to that, though, I want to just bring out more clearly the contrast between measurement-omission and context-omission. In context-omission, you divide different things into different elements, and focus only on one subset of the elements. Therefore, after the process, you have *less* in view than you did before. In measurement-omission, you interrelate the different things, so that after the process, you have *more* in view than you did before, only you have it in view in a less taxing form—in a more economical form, in the form of one line rather than many dots. After context-omission, one is aware of a unity that was there before one performed the process. For example, if we we imagine that you could perform context-omission on birds, what you'd reach is beak and feathers, and beak and feathers was there in all the different birds before you performed the process of context-omission. So you become aware of a unity that already existed. The process of context-omission is just the process of clearing away distractions. Measurement-omission is the *creation* of something. It's the creation of a unitary way of viewing the multiplicity that was there all along. A way that's based on facts about the multiplicity, on the relations in which they stand, but yet is still something new—not a new object, but a new *way* of being aware of the objects that you were aware of before. There are other, related contrasts that we'll get to later when we bring in the other points important about concept formation and measurement omission.

2.2.8. Multi-dimensional atributes

Before we turn to those, I want to make one more point about measurement-omission as such, and particularly about conceptual common denominators. They needn't be one dimensional. We've been talking about one-dimensional CCDs—things that can be represented on a chart by a line—things like hue or length or loudness. But CCDs can be multidimensional. In fact, colors vary from one another along at least three axes: hue, saturation and brightness, and most of our color concepts

subsume three dimsnsional ranges. If you can build a 3-D graph, which you can in some computer programs, or build a model, the concept "pink" would have to be represented by a kind of cube or other 3D shape, rather than by a line. Pink's a range of brightnesses, a range of saturations, a range of hues. In fact, many CCDs are much more complicated than this. Look now at our little Exhibit A. I've drawn a bunch of things that I think are recognizably man-shapes—I mean, they're simplified, but you can recognize them as looking like a man, and thre's one that looks like, I don't know, a dog or some other four leged animal. I did this in a computer program which made it easy to figure out how many numbers you'd need to specify the differences between these things: they differ from one another along thirty-five dimensions. That is, if you had some kind of a thirty-five dimensional graphing program, you could draw a "box" (you know a... whatever thirty-five dimensional shape would be) that would enclose all the men, and not the dog-shaped thing. So you can see that even such complex shapes can ultimately be reduced back to linier measurements of the sort we've been discussign. Now granted, actual human shapes are more complicated than these guys here, but you get the point. And notice that you can see that the men look similar to one another without constructing those thirty-five dimensions. And if I asked you, "how are they similar?" you'd say, "well, the legs are longer or shorter, they're in this or that orientation," and, probably, you could in this way eventually list the thirty-five axes. (Each man is made up of a number of shapes. One for the head, one for the torso, two arms, and two legs, a neck, etc. Each of those shapes is in some relative horizontal and vertical position relative to the others, each has some thickness relative to its width. In the case of the oval that makes up the head, it's got some diameter, etc. They're those kind of basic measurements for each of the shapes, and then measurements for their orientations relative to each other.)

2.3. Concepts of Entitites

Okay, I want to turn now to concepts of entities. We were focused before, in explaining measurement-omission, on concepts of single attributes, and, in particular on one-dimensional attributes, because that made it simpler to explain measurement-omission and to see how different it is from context-omission. But the features that made these examples particularly good for this purpose are atypical, and they obscure some important points about the concept-forming process. We can get a more complete picture of the process by thinking about concepts of entities. The factors that are going to be present in concepts of entities, but absent from these concepts, like "length" or "blue," are present in most concepts, so the entities are really the case to focus on.

2.3.1. How entities are measured and integrated

Now, the first point is that entities are measured by measuring their attributes. For example, I'm measured by measuring how tall I am, what I weigh, the relative lengths and widths of my parts, etc. And a group of similar entities is formed by differentiation from other entities along some CCD, often shape, though it needn't be shape, and in more complex cases, won't be. You can also measure my intelligence, the range of my awareness, my metabolism in various ways—things other than my shape. But, in general, a group of similar entities is formed by differentiation from other entities, the foils, along some CCD shared by all the entities, and often, at least in the simplest perceptual cases, the CCD is shape.

The units of one of these groups are similar in that they share what Ayn Rand called a "distinguishing characteristic." And this distinguishing characteristic is a range along a CCD. So for men it might be the characteristic "man-shape". Now, men share a lot of distinguishing characteristics, but the distinguishing characteristic we share along the CCD shape is man-shape. And that's the thing that's, in a simple way, illustrated (in Exhibit A) by these man-shapes as opposed to this dog-shape,

which is along the complex thirty-five dimensional CCD that I set up. And, of course, it would be a more complex measurement in the case of actual human shapes, which are more complex than those.

When one forms the concept of an entity, one forms it by omitting measurements, but it's not just the measurements within the range of the distinguishing characteristic along the CCD; it's not just the range of man-shapes that one omits. Rather, when one forms the concept of man, one omits *all* measurements along *all* axes. One omits the measurements of all the known characteristics of man, for example, perhaps, the way when men move, the sounds men make, the texture of their skin, their coloration, etc. (and notice that men are similar to one another in all of these ways as opposed to other existents) and one establishes a policy of omitting newly discovered differences between men: "When I discover new features of men, I'm going to omit the measurements of them in regarding them as men". Thus one will omit measurements of their intelligence, their occupation, their marital status, their political affiliations, and so forth. A concept of an entity refers to those entities possessing this distinguishing characteristic along this CCD, with whatever other characteristics these entities may have, in whatever other degrees they many have them, along whatever other axes they may be measured.

Thus, Ayn Rand's definition of a concept. A concept is "a mental integration of two of more units possessing the same distinguishing characteristic(s), with their particular measurements omitted." Notice she says they have the same distinguishing characteristics, but the particular measurements that are omitted are not *just* the measurements of the distinguishing characteristics, but all the measurements of the existents. The particular, distinguishing measurements of any particular existent within the category is omitted. Now this definition applies not just to concepts of entities, but to concepts for any existent.

2.3.2 Generalization of the preceeding point to all existents

"Existent" means anything that exists, be it an entity, and attribute, a relation, an action, etc. Whatever existents one conceptualizes, the concept is an integration of the whole existents, with all their measurements, within relevant ranges, included. And our first examples, the examples like blue and length, were atypical in that they were examples of concepts for singular characteristics, so that there are no other measurements to omit. But if you form a concept for an existent that's complex enough to have different attributes which can be measured along different axes, then the concept for it includes the ranges that the units fall into along each of those axes. You distinguish the units, say, men, from other things along a certain axis, along a certain CCD, say its shape, at first, but then the concept refers to all the things that fall within that range along the shape axis with all their features, whatever they may be, in whatever ranges they may be, along other axes. So the CCD and the distincuishing characteristic (DC) are special in that they're the means of distinguishing the existents being conceptualized from other existents. But they're not the only ones whose measurements are omitted.

Here I think the contrast with context-omission is clear. In context omission, the thing that you focus on—the distinguishing characteristic, now understood not as a range but as an identical thing in each of the units—is all that you focus on, and you ignore, you omit, all the characteristics in virtue of which the units differ. This yields only a partial awareness of the units, with their differing traits excluded. So we'd have to exclude from the referents of the concept "man" all the traits in which men differ. In fact, there are *no* traits in which men don't differ, but suppose that there were some such trait—say, rationality, and we can just keep that constant. Well, then, on the Moderate Realist, context-omission view, "man" would have to refer just to the rationality and not to the other traits. By contrast, measurement-omission doesn't get rid of difference, it interrelates differing things, and this interrelation can include elements, characteristics, other than those that are the basis for the interrelation. (The one that's the basis for the interrelation is the distinguishing characteristic along the

CCD.)

Now, I'm not going to stop here to talk at length about the analytic-synthetic dichotomy, but (for those of you who have read Dr. Peikoff's article on it), I want to point out that this is where Ayn Rand's theory of concepts gets rid of the analytic-synthetic dichotomy, and it's not getting this point that makes all other theories of concepts prone to it in one way or another. It's precisely insofar as all prior theories of concepts fall back on accepting, some way or another, context-omission, because they don't have a view of their own, that they end up accepting the analytic-synthetic dichotomy. And again, context-omission is, in a sense, peculiar to Moderate Realism—I think it's the essential process of Moderate Realism—but since it's the only process that's been defined for the forming of concepts, it tends to get taken over by other views, in certain situations, even though that's not their official line (I can give examples of this later if people want). And it's because of this that they fall into the analytic-synthetic dichotomy. People who want to hear more about that might consider taking Ben Bayer's class where he's talking about the dichotomy and about some modern criticisms of it which are not any good.

2.3.3 How essential similarities make possible the integration whole entities rather than merely their distinguishing characteristics

Now, I've said that the concept "man" integrates the whole entities, men, not just their distinguishing characteristics. But how does it manage to do this? How is it that by differentiating men according to one characteristic—his shape, let's say—how does that enable us to, by omitting the measurements of that characteristic, form a concept for the whole entity? Why don't we just get a concept "man-shape"? If the question's a little unclear, I think it won't be when I give this parallel example.

Consider two concepts: "man" and "greenie". A "greenie" is any existent that is green—any green entity. It doesn't just refer to the green color of the entity, but to the whole entity. It refers to "those things whose color falls within the green range, along with whatever other traits they might happen to have, in whatever ranges they have them"—like the concept "man" refers to all those entities whose shape falls within the man range, along with other traits in whatever ranges men have. Well, why isn't "greenie" a good concept? Why don't we have a concept "greenie"? Betsy?

BS [faintly]: It's not cognitively useful.

GS: Why isn't it cognitively useful?

BS: It doesn't put together things that you need to inter... uh, to act with together.

GS: Well, why don't we need to act with "greenies" together? I mean, I think all that's right, but: "greenies" have nothing in common, other than being green. I made a little graph, it's a little silly, but, if you look at Exhibit B, you know, I just picked something random to chart against hue in this graph: size. If you plot entities, on a graph according to their hue and size, you'll get a random bunch of dots. Green things are not about the same size as one another. However, man-shaped things have about the same range of awareness as one another. If you plot "shape of man" against "range of awareness," you'll get, a reagion: the things that have about the same shape have about the same range of awareness. The things that are monkey-shaped have about the monkey range of awareness. The things are green-colored don't have green size. There are green eyes, green cars, green buildings—no correlation. Men share many characteristics, and if you tried to chart any two characteristics against each other, most of them, you'd get boxes for man, whereas you'd get nothing for "greenies," you'd get this kind of scatter graph with a bunch of dots. These this charts are a little silly, but I think they make a certain point.

Most notably, men act similarly, as opposed to animals. They move about in a distinctive way; they make sounds in a distinctive way, etc. Now why is this? Why is it that being a man (or having the

man-shape) correlates with many other traits, whereas greenness doesn't? Well, because man's shape is causally connected to all these other traits, whereas greenness isn't. In a way, man-shape causes the way men move. Now that's not the whole story, and causality is complicated, and in another way the way men move causes the shape because there's an evolutionary process. But what's important here is that there is a causal connection, and we can capture it, at the most basic level, by observing that is a way in which my shape causes the way I move. I move like *this* because I'm shaped like *this*. If I was a big cube I couldn't do that—I couldn't walk. I know: I had a period in my life, I don't like to talk about it. [Laughter] Similarly with man's sounds: my speach comes out of my mouth. If I had a parrot beak... Well, parrots can make similar sounds, but if I had a really different-shaped vocal apparatus, I couldn't sound like I do (and it would be disappointing). Notice also, that even within the man-range, shape affects motion. Think of John Candy and John Cleese: differently shaped men, they move differently. John, you know, Cleese is the tall, thin—he used to be, anyway—guy in Monty Python, he could do those silly walks; John Candy was a big fat round guy, and he couldn't have done those silly walks.

It is this sharing of many interrelated characteristics that make men commensurable as wholes, as opposed to merely having some of their attributes commensurable. A green leaf, a green car, a pair of green eyes share only one attribute. The attribute is commensurable, but the entities are not. There is a holistic commensurability between men because they share many causally connected attributes, all of which fall within ranges. And this is what makes the unit perspective on men as such, as opposed to just on one of man's characteristics, possible and valuable.

Men are alike in a complex way that makes all of one's knowledge of any given man relevant to every other man. Men can be compared and measured against each other holistically, not just in terms of an isolated part. You need this for concept-formation. You need the existents that you're forming a concept of to be distinguishable from other existents along some CCD, but you also need the distinguishing characteristic along that CCD to relate to many other characteristics, and to cause or stand behind many other characteristics in such a way that the entities are comparable to one another as wholes, not just that they have one attribute comparable.

Another way to put this point is that concept formation requires an essential similarity: one from which follows many other similarities. And as we'll see in more detail tomorrow, I think this is even true for first-level concepts: ones formed directly from perception. But more on that tomorrow. Concept formation requires an essential similarity, one from which follows many other similarities. Now I just want to stress the importance of this, because you can find *some* similarities between anything. Here's a particularly gruesome example I came up with. Take the CCD "importance attributed to helping the least competent people." You can measure and compate people in this respect—along this axis. For example, Mother Theresa would be all the way on one side of it, right? Ayn Rand would be all the way on the other side of it. Franklin Roosevelt would be somewhere in the middle, but much closer to Mother Theresa. Where would Adolf Hitler fall? He'd fall sort of on the Ayn Rand side of it, because he thought we should crush the least competent people in order to make the race as a whole come out better. If you think of this as your CCD, and you tried to form a concept for the whole people, you know, not just for some trait (and here there shouldn't even be a concept for the trait), you'd wind up with a concept, maybe you'd call it "Facist", under which you'd group together Ayn Rand and Hitler. And then, for example, if you started using this as a basis to compare them as a wholes, you might read Atlas Shrugged and find gas chambers and things between the lines of it, as some terrible people do. This is what happens when you group by nonessentials. You destroy your mind.

Now that's an example from a high-level concept, and there's a whole lot of mistakes that you would have to make, or lies that you'd have to purposely be telling yourself and others, in order to get into a position where you can make that kind of error. But it is an evil version of—and a possible version of—a concept like "greenie". "Greenie" was supposed to be a first-level concept and it's just

not possible to go wrong like that on the first level. I think this Facist example indicates how many premises are behind our understanding of what is similar and how it one similar trait relates to another in complex cases. And that's going to be very relevant to the topic of hierarchy, which, again, we'll discuss later—in this case, on the last day.

2.4. Definitions

Now I've been talking about essences. I said a concept requires an essential similarity. Ayn Rand discusses essences, of course, in connection with the topic of definition in ITOE. And let me end today's lecture by saying a little bit about definition, and then we'll take questions. I'm going to talk about the purpose of definitions later, when I talk about the purpose of objectivity. For now, all I want to say about about what they are is this: a definition is a statement of the way in which one is differentiating the units of a concept from other existents.

We said that to have a concept is to be on a policy of regarding things a certain way. A definition is a statement of policy. Now, definitions are given in terms of other concepts. So a child can't have definitions for his first concepts at the time when he first forms them—not what we normally think of as definitions: verbal formulas stating the policy—because, by definition, if they're his first concepts, he then doesn't have other concepts in terms of which he can state his policy. But what a child does is he resolves, in non-conceptual form to differentiate on a certain policy. (Remember there are actions of our consciousness which we can do and have their content be perceptual or conceptual.) And we can articulate the policy on which he's differentiating, and the articulation of the policy would be a definition. So, although the child doesn't have a definition himself, he's acting implicitly on a policy, the articulation of which policy, in conceptual terms, would be a definition.

2.4.1 The parts of a definition

Now, how does one articulate a definition? A definition has two parts: the genus and the defferentia. Genus is just Latin for "kind" (and also Greek for kind, really). The genus in a definition specifies a wider class of commensurable things within which the units of the defined concept will be differentiated. So, for example, if we're going to define man, animal would be his genus, because animals are commensurable with men, and we differentiate men from other animals. The differentia, which is just Latin for "difference," is the part of a definition which shows how the units of the concept are different from the other things within their genus. So, for the young child, the differentia of man might be "man-shaped," which he would hold perceptually. For us, it's "rational." Man is the rational animal. The differentia is the distinguishing characteristic, and it must be the fundamental distinguishing characteristic: the one that causes all the other known distinguishing characteristics. It's the essential characteristic.

What we said an essential characteristic was, was one from which many others follow. That's not quite accurate, I think. Fundamentality is the name for the fact that many things follow from something. Essentiality is that same fact viewed from the perspective of the cognitive use which we make of it. So the differentia in a definition is a fundamental distinguishing characteristic. It's the characteristic, of those that differentiate the units from other things, which causes all or most of the rest. That's what it is to be fundamental, that's essential relevant to our need to use such a characteristic in order to get a handle on the things, in order to think about them.

2.4.2. How definitions can change

One last point under the heading of definitions: when one conceptualizes a group of essentially

similar existents—that is, when one groups of existents that have fundamental characteristics in common—there will be a lot to learn about them. When one conceptualizes a group of essentially similar existents, there will be a lot to learn about these existents. And as one learns more, one's definitions can become obsolete. There are two ways this can happen: you can either discover a new characteristic of the things you're grouping which is more fundamental than the one that you had previously thought to be fundamental—that one *was* fundamental to all the ones you knew, but now you've discovered one that's fundamental even to that one—, or you can learn of existents that have the old differentia but which nonetheless are importantly different from the other units of the concept. Say you encounter a man-shaped and man-sized puppet or a manikin, in the stage where your implicit definition of man is "man-shaped thing." Well, the manikin's a man-shaped entity, but he sure isn't like other men. So he's *not* a man. And your old definition is obsolete.

The essence is the characteristic that distinguishes the units from all other units *within the context of your knowledge*. And as that context changes, the essence will need to change. But when the essence changes, the concept, the integration, the mental entity remains the same. Why? Because you're still integrating the same range of units, and you're still doing it on the same basis. The initial basis for the differentiation and integration of men was *not* "man-shape" alone, in splendid isolation from men's other characteristics. Rather, you were differentiating and integrating men as wholes, based on a whole set of similarities and differences from other things, of which you saw man-shape as the fundamental (if you did see man-shape as the fundamental). You saw the set of similarities and differences as centered around the essential characteristic, and you used this to get a handle, to get a grasp on what is, in fact, and what you already knew to be, a multi-aspectual similarity between the units. Now you continue, after you have the new essence, to differentiate the units globally, you've just got a kind of new grip on how you're doing that.

2.5. Concluding remarks

What we did today is we looked at what a concept is, and what the act of measurement omission, by which one forms a concept, is. Notice we did not discuss norms or how to do it, and we weren't really distinguishing between levels of concepts very much—that is, between concepts formed directly from perception and ones formed later. We're going to do both of those things over the next two days. Tomorrow we're going to discuss the process of conceptualizing as a whole. We're still not going to be focused on norms and guidance, but we're going to be focused on the whole process, of which concept formation is the juggernaut, the spark, the main impetus. We're going to talk about on what basis you initially form a concept, how you then apply that concept to objects in the world, and how you go on to form higher and higher concepts. All of that, tomorrow. That's why the heading for tomorrow as a whole is "conceptualizing", whereas the heading for today was "concepts". Then, on the last day, we're going to look at that whole process from another perspective, from the perspective of the fact that this is a process that you have to enact and sustain: that it's a volitional process. And we're going to draw out methodological guidance from it. In fact, all we're going to do is just restate what we said in the previous days, but now it's: "and you've gotta *do* this, as opposed to these other other things, which you might try to get away with—and you can't."

Q&A

Um...go ahead, Don.

Don[clearly audible]: Uh, well... what is the validation of this theory?

GS: The theory of concept formation? It's the observation of... the introspective observation of what you do in forming a concept, and what you need to do to form the concept. That is, we can

distinguish two things: “what is the validation of a given concept?”, right, and “what is the validation of this theory of concepts and of their validation?” We’re going to talk about how one validates a concept tomorrow and then over the next day. And we’re going to talk about how one knows that that’s how you validate a concept in the context of doing that. The validation of the whole theory is just the sum of the fact that all of our concepts have to be validated in this way.

Let me, sort of, just point out a certain possible error here. It’s a mistake to look at concept formation as though you were a psychologist. That is, it’s okay in some senses--if you’re a psychologist, but you’ve got to be careful--it’s not that we just have a whole bunch of concepts and then we study how we form them. This isn’t primarily a topic for cognitive psychology. Compare what I call “the problem of concepts” to what you might call the “problem of percepts,” that is, “how is it that you perceive? How is it that you manage to integrate a whole bunch of data impinging on your sense organs into an awareness of 3-D moving entities?” That’s a perfectly valid, sensible question, but it’s no part of philosophy. That is, it’s a question for people studying your brain and your neurons and... and it doesn’t matter--I mean, it matters in the way all knowledge matters--but it doesn’t matter for you as a perceiver. The best scientist of perception--Gibson, or whomever, or someone in the future--perceives no better than anybody else. Right? It doesn’t help him, at all, perceive. Now, whereas Ayn Rand, the best theorist of concepts ever, conceived a lot better than everybody else. And we, as we learn this theory and learn to integrate it into our lives, learn to think better. Now, incidentally, that’s not the emphasis of this course: learning how to think better. It’s “how to understand the structure of this theory.” But if you do come to understand it and internalize it, it’ll improve your thinking.

Now why do I bring this up in answer to your question? Well, if you take the psychologist model of it, you take what concepts there are--what words people use and how they use them--as a given, and you think that what you’re doing is studying how it is that they go about doing that. And if you do that, you’ll never come to this theory, because, especially with higher-level concepts, people *don’t* do that, and that’s why they’re so confused. What you have to do is think about it from the perspective of somebody who knows how to conceptualize--implicitly, first--who’s able to think, and then you think about “what am I doing in thinking, and how does that help me get at the world in the case where this is... where I’m able to and where I know how I’m able to?” And then when you identify how you do that, you can apply it to more and more complex cases. And that’s what Ayn Rand did; she found this theory by asking herself, “what is it that my mind does when I form and use concepts.” And it’s the same thing that Aristotle did in discovering the laws of logic. That is, he didn’t go around surveying people and seeing how they argued. Not that you were suggesting that, but I just... that’s a thing people often try to do, and I think that that has to be differentiated.

Okay, I don’t know your name, but you.

John[faint-can be heard clearly by turning up the volume, but then your inserted comments are too loud]: My name’s John.

GS: John? Go ahead.

John: You mentioned that when you form a concept you [scraping noises in foreground] omit measurements--not only measurements that you’re aware of...

GS: Mm-hmm

John: ...but a whole host of measurements that you may not be aware of yet.

GS: Right

John: I was wondering if, in that process, any of the measurements omitted might be non-commensurable...

GS: Um... okay, he says, “you, when you...”... I said that when you form a concept, you omit measurements--not only the measurements you’re aware of, but also, um, other measurements along unknown ranges, uh... unknown axes. Might those be... might any of those other measurements be non-

commensurable? No, because if they were, they wouldn't be measurements. That is, a measurement, uh, is um... "a measurement is a measurement," I wanna say. What one measures is, to speak metaphorically in terms of axes again, which I find very helpful--often there are things that one can't... one needs to just make the metaphors into concepts, one can't form a concept independent of the metaphor, uh, which then would justify the metaphor because, uh, to do that you'd have to have all your concepts before you started, um, so to speak in terms of axes: what a measurement measures is something's position along an axis. And it's gotta be an axis along which other things have positions. So, um, you can't have a measurement that's not commensurable with anything. Um, now, um, so if you're going to form the concept man from, say, uh, me and Don as opposed to the light fixture, you know, uh, the... you'll omit each of my measurements. Each thing that's measurable in me will have some correlate in him, uh, if it d-... I mean there might be some that don't, but they're... they would be handleable in some other way, and it wouldn't... it doesn't... you could practice something like context omission on them, if that were the case. You could say the fact that, um, Don is, uh, wearing an OCON badge, and I'm not, say, is no part of the concept "man," except that... even that it is, because it wouldn't...it would...it, it doesn't fall under, you know, part of his shape, the OCON badge, but it forms under, uh... falls under "action he's performing", or "location of things re-re-relative to him, and, and so forth, all of which is measurable. Um...so I, I don't think that's a possibility, I'm not even sure what it would mean for it to be a possibility.

Um...uh, Betsy. Betsy was....

BS[clearly audible]: Okay. [clears throat loudly] One thing I have a hard time introspecting about, or finding in Ayn Rand, is, uh, the idea of *acts* of measurement omission or a process of measurement omission. It seems... it seems as if what she's saying is that we are grouping by similarity, and when we do that, we omit the measurements, like it's a *result* of doing something, rather than something we're doing.

GS: Well, I do think she uses the word "process," um, so I mean, we'd have to... I don't a... the quote next to me. But I think the, the... we can rephrase the question to, "what is it that happens consciously?"

BS: Right

GS: Um, to what extent is this, um, a conscious process one performs, and to what extent is it, in effect, the subconic-... -conscious mechanics of the process. And this comes up in the question period, uh, in the workshops that are transcribed in the current edition of ITOE; I think Professor B asks this question. Um...measuring... in a sense it's the subconscious mechanics, and what one does is observe similarities. But, uh, I think more of it is conscious than that phrase might suggest. Because what one has to do is first observe similarities between things. Now what is it to observe similarities? Well, as opposed to what? As opposed to not observing any relationship between them at all, right? Um, as opposed to not distinguishing them from one another in the first place, because only if one sees them as two things can one see them as standing in a relationship to each other. So we... as opposed to disparateness, as opposed to a lack of distinguishing, and then as opposed to distinguishing but seeing them as different. Right? So those are three things. So one has to, um, see them as similar. And, uh, and what it is to see them as similar is to see them as less different from one another than they are...it, it... something else. Now, that puts words to it, and a child first doing this wouldn't have words to it, but that does aptly name the content of what you experience when you just look at two shades of blue next to a shade of yellow. "Oh, those go together as opposed to that. They're closer than that." So that, I think, *is* in consciousness. And then what you do is you, um, see another shade of blue: "Oh, this goes along with these," you know, um...the process of integrating to form the concept is a process of... I think there's something conscious that corresponds to omitting the measurements, and it's the... the... not just...but... not just "this one and that one and this one, but any one like that." And I think that is

conscious. I know I do it all the time, when I form concepts, and I think that that's what a child does, and that's what it is to omit measurements.

BS: ...Well, is that similarity... similarity grouping, and as a consequence the individual measurements are omitted?

GS: No, the similarity grouping is sim-... similarity is a three-place relation, right? "This is like that, as opposed to that. And this is like that and that, as opposed to that." And... it's something one does to things sort of individually to put them in a group. What one does to get beyond that is to say, "Look, no more of this 'this and this as opposed to this'. *All* these as opposed to those." And there is an... I think, an introspectable act of doing that. One doesn't think of it as "omitting measurements" (unless one does because one's read up on this theory) uh...because one doesn't think of oneself as measuring. But what it is to notice similarity is to notice a kind of relationship between things, and that is, in essence, measuring. Um, so that's... that's how I would characterize the conscious to the non-conscious aspects of it.

Um... Guy, is it?

GB: Yeah. Um... I... like, I have a question about concept formation

GS: Actually, could I pause for a minute? Matt, are any of these voices from the people getting on the thing? They are? Okay, good.

GB: I have a question about concept-formation as it relates to axioms. I just want to know...

GS: okay

GB: ...um...if you're going to cover that later, and then I won't even ask.

GS: Ask it, [laughter begins] and then if I'm going to cover it later... let... let's do it that way.

GB: Okay, uh, the question is like I... like in... normally when you're forming concepts, you have to differentiate it against something...

GS: Yes

GB: ...like you differentiate dog from the people,

GS: Uh-huh

GB: but in ax... axiomatic concepts, um, you can't really do that, right? You can't...

GS: Yes

GB: Like, you can't... what are you going to differentiate consciousness from? Or existence from, or something?

GS: You don't. Um, the question is, "In normal cases of forming concepts, there's a differentiation, there's a foil, uh, something that you differentiate the units of the concept from. Not in the case of axiomatic concepts. Uh, I think in a way that's right. You don't form the concept "existence" by finding some non-existence to contrast it with, there aren't any. Um...but you also don't form the concept "existence", as an explicit concept, until quite late. And, um, it's a kind of special case. There are other kinds of special cases, like what do you contrast length with to form the concept "length"? Well you contrast it with width, I suppose, but width isn't something that varies from length along a CCD. At least, I don't think it is. Um... I think the problem, the question that gives rise to the need for that aspect of this theory, uh, that is, the measurement omission along a CCD where there's a contrast with foils, uh, doesn't arise in certain cases because they're sort of too simple cases; they have certain special features. Uh, those... concepts for those things, I think, can only be formed when one's in a context of thinking in concepts, so one has a lot of concepts already, and one is not trying to bring into that mode of being aware--uh, that is, conceptual cognition--a new sort of thing. Uh, so there is a kind of contrast that takes place, I think, in the forming of any of these concepts, but it's not a normal contrast along a CCD. It's often a kind of psychological contrast between what you do with your kind of awareness, uh, and, um, with one ki-... with awareness of one attribute: that you do with awareness of certain things or certain perspectives on things and what you do with awareness of other things and

other perspectives on things. In the case of axiomatic concepts, the issue of psychological time-measurements is important, and that's a little esoteric-seeming, um, but, um, what I think you do is, uh, recognize that your wareness of certain facts plays, and has to play, a certain role in your consciousness, and differentiated from the role played by other facts. Um, and therefore you recognize a need to hold these facts in a certain way. Uh, and there is kind of measurement-omission involved in the steps of that process. For example, the axioms are things that are present in all cognition, rather than just some, and which need to *always* be held in some form, as opposed to *sometimes* be held (you see, more, less) but what you're not omitting is measurements of the things. Or rather, you are omitting measurements of the things in the case of existence--you're omitting all measurements of all things. But what you're not differentiating is those from other measurements of, of other things (which, by definition, don't exist). But that's uh, uh, hard case, and I'm not... I'm gonna talk about what axiomatic concepts do for us a little bit later, and I'm going to talk about the fact that their basis is in perception, but I'm not going to talk about that, so... now I have, and so thanks for giving me the opportunity to.

Um... any other questions? Uh, Ben.

BB: Two questions about the theory of context-omission. One, you mentioned a scholar who gave it that name: who is the scholar? And two, uh, you said that other theories--other false theories--of concepts sometimes fall into that explanation.

GS: Okay

BB: Could you give an example of that?

GS: So the question is, "I referred to the theory of context omission as the "realist theory of concept formation," and I said that I got that phrase, "context omission" from a scholar, and could I say who it was? And two, could I give examples of other, non-Moderate Realist theories that use context omission somehow?" The scholar's a guy named John Linnell, who was a student of, uh, Wilfred Sellers, who was a sort of major figure in um, in uh, (or a fairly large figure anyway) in analytic philosophy. Um, and I suspect that the insight is as much Sellers' as it is Linnell's, um, but, uh, 'cause the perspective that Linnell takes in this paper, which is useful, is a Sellers-y kind of perspective. Um, but the guy is Linnell, and it's in a paper called something like, "Locke's theory of abstraction." [In fact, it is titled, "Locke's abstract ideas."] The guy only ever wrote two things, so if you look up, in *Philosopher's Index*, it's one of the only two that'll come up, and you'll be able to find it. Um...Why I um...I mean, he used, I think, it kind of as a kind of in passing as a... a phrase, but I really liked it, uh, I think he, or somebody else, also called it a "screening operation," which is a nice metaphor.

Um, now other theories that use this. Well, first of all, um...most Conceptualist theories, or many of them, uh, tend to be realist about certain concepts and Conceptualist about others. Uh, Locke certainly is. Locke is a realist about concepts of simple ideas, concepts basically for sensations, and he is a Conceptualist about other concepts. His conceptualism consists in the view that when you form concepts for things other than sensations, um, the arbitrary element that comes from your mind is just grouping together concepts for sensations. So any concept, for Locke, is going to be formable by context omission, um, because its units are always going to have in common the basic sensory qualities that had to be grouped together to get to that concept. Um, but he's aconceptualist rather than a realist, because he thinks that, without a certain mental act that you perform, there would be a kind of, just, dis-unitary heap, rather than a unitary cognition. The heap wouldn't be a heap of different units, but it would be a heap of different parts of each unit. So the view is in Locke, and I think in the same way that it's in Locke, it's in a lot of empiricist sorts, um, but Locke's an example. It's also in Kant, if you read Kant on concept formation in, um... or on concepts, mostly in his lectures on logic, he sounds just like Locke, um... I mean, it's really shocking. The difference is that he thinks that the percepts from which one abstracts the um... you know, things to group together, like Locke groups them together, um,

the sensory bits that then need to be ordered, uh, are already put together, by, in effect, implicit, innate concepts that one already has, uh, the... the so-called categories, and more than the categories... there's a whole apparatus in Kant. Um, but Kant holds this, and Kant has really strikingly good descriptions of this. He says things like (this isn't an exact quote, but it's a quote), um: "a concept is a representation of an intuition (which is his word for percepts) with parts of it taken out." It's a partial representation of an intuition. Things that sound really, um, uh... it really struck me over the head because it sounds almost word for word, with some substitutions, to um, to phrases Dr. Binswanger has used to describe this view. Um... so that's an example: Locke and Kant. Um...I think there are others, but, um... none are sort of rushing to mind at the moment, but that should be enough to give the sense that it's in a lot of places. Oh! Also, one more. The, the uh, in psy...in psychology, and particularly psychology commenting on, um, philosophical theories of concepts, one often encounters the phrase, "classical model," or "classical theory" of a concept, which holds that a concept is like a bunch of words in your head, like a definition. And people who hold that, hold that th...I... generally hold that the, you know, individual words that are parts of the concept or... refer to parts of the thing, and it amounts to the same theory. And Locke is often given as an example of this. Also Plato and Socrates are given examples of it, and they are, but they're... Socrates, at least, was a Moderate Realist, and Plato, in certain moods, was, but usually an Extreme Realist.

Lecture 3: Conceptualization

I'd like to begin today with a somewhat lengthy quote from Ayn Rand, *The Virtue of Selfishness*, "The Objectivist Ethics," page 21-22:

> The process of concept formation does not consist merely of grasping a few simple abstractions, such as "chair," "table," "hot," "cold," and of learning to speak. It consists of a method of using one's consciousness, best designated by the term "conceptualizing." It is not a passive state of registering random impressions. It is an actively sustained process of identifying one's impressions in conceptual terms, of integrating every event and every observation into a conceptual context, of grasping relationships, differences, similarities in one's perceptual material and of abstracting them into new concepts, of drawing inferences, of making deductions, of reaching conclusions, of asking new questions and discovering new answers and expanding one's knowledge into an ever-growing sum. The faculty that directs this process, the faculty that works by means of concepts, is: *reason*.

Yesterday, we discussed, in effect, the process of concept formation, but thinking of it more or less as something that ends with a simple abstractions as "chair," "table," "hot," "cold," etc.—or, at any event, we didn't discuss it insofar as it involves going beyond such simple abstractions. Today we're going to look at the process as a whole: the process of conceptualizing, which needs to be actively sustained. But we're not going to look at it from the perspective of its needing to be actively sustained, and what one needs to do to sustain it. That's the topic of objectivity and of logic, which is our subject for tomorrow. The topic for today is the process as a whole, but viewed not with the special emphasis on what we need to do as directors of our own thinking; it's the nature of the process of conceptualizing taken as a whole.

Let's begin by looking again at the formation of concepts. Concepts are formed based either on earlier concepts or directly on perceptual evidence. And so we distinguish between first-level concepts—ones formed directly from perception—and higher level concepts, or abstractions from abstractions—the ones that are formed and require previous concepts for their formation. Let's look first at the process of conceptualizing on the first-level: the formation and use of concepts which can be formed directly from perception and then applied directly on the basis of perceptual knowledge.

3.1. First-Level Conceptualization

3.1.1. How first-level concepts are fofmed directly from perception

First-level concepts are formed directly from perception. They integrate the things which are differentiated from one another in perception, and those things are, as we've already noted, entities. A concept is based on an essential similarity between the existents conceptualized, between the units. And so in order to form a first-level concept, one must be aware of an essential similarity between the entities that one perceives, or between some of the entities that one perceives, and we can be aware of this perceptually. That is, we can perceptually grasp similarities and differences between things. We can perceptually grasp the similarity of man-shapes—that is, the shapes of different men—as against dog-shapes, the similarity of a man's motion as against a dog's motion, the similarity of the sounds a man makes as against the sounds a dog makes, and so forth. And as a whole, we can grasp, perceptually, the similarity between men as against dogs.

But moreover, we can grasp the causal relations between the different attributes of man already

at the perceptual level—not *all* the causal connections between *all* his attributes, but many of the connections between many of his attributes. We can grasp that man-shape causes man-motion. Now that's not the full story, again, about the causing of man-motion, but there is something that we grasp that I'm articulating by saying that. Namely, that I can walk this way because I'm shaped this way, whereas were I a big cube, or a ball, or a dog, I wouldn't be able to move in quite that way owing to my different shape. (If people are interested, we can talk in the question period about the different sorts of causal relations there are, and which ones can be grasped in which ways. I think that's an interesting topic, but it's not one for today.)

The point I want to make now is that it's not that we perceive the man-shape and the man-motion as separate things, and then, later, you know, as a separate part of the perception, perceive their relationship. No, rather we perceive the shape and the motion as a unity. To perceive a man is to perceive an entity acting over time: an entity that moves through the world in a certain way that's inseparable from its shape. So it's not that we perceive all these different bits, or things, that we would judge in all these different discrete propositions: that guy's got a certain shape, and he walks. Rather, we perceive a whole thing, of which these different attributes that we might talk about separately, or actions that we might talk about separately, are abstractions. It's a mental act to focus on just the shape or just the motion. One perceives the thing as a whole. Look at that fly that's flying across the room: you just see the moving fly, you don't see a black dot here, and then one here, and something in between, and little wings; you just see, you know, a fly. We grasp the attributes and causal relations at the perceptual level as part of perceiving the entities. That's part of what it is to perceive an entity rather than a sundry collection of sensations.

Now, none of this knowledge of the different attributes, of their relation, etc., is in conceptual form, and it doesn't need to be conceptual form for us to use it to form the concept. Rather, it's held perceptually. And remember, even on the first day when we talked about the things an animal can do based on its perception, such as have desires and so forth, we were making the point that you can do things with perceptual knowledge. You can perform mental acts whose content is perceptual knowledge. Given what we've said yesterday, I think that's all we need to say now about how one can form concepts directly based on perception. To form concepts of entities, you need to see that the entities are similar, and you need to see that their similarity is essential—that some similarities give rise to many others, and you can see all that in perception.

3.1.2. How first-level judgments are based directly on perception

So I want to move on, now, to the question of first-level judgments. That is, judgments which are based directly on perception. Now, what is a judgment? A judgment is a conceptual identification of something. To make a judgment is to identify something under a concept. And the form of a judgment is a proposition. To judge is the act of subsuming under a under a concept, and the form of the content of a subsumption under a concept, is a proposition. A proposition has the form "S is P," where S and P, subject and predicate, can be concepts or else descriptions made by compounding concepts. In a judgment, you're identifying the subject, "S," by subsuming it, or some characteristic of it, under the the concept "P." Now, there are mental states other than judgment (I just mention this incidentally) that also have propositional form. For example, you can hope that S is P. You can hope that Objectivism is winning.

In a judgment, the subject of the proposition is identified by subsuming it, or one of its characteristics, under the predicate concept. The basic case of this is some sentence like, "That's a dog." And "That's a dog" is exactly equivalent to a child going, "Dog! Dog! Dog!" because the word "that's" is just a grammatical device for integrating a direct perceptual awareness into a sentence. In both cases—"That's a dog," and "Dog!"—the subject is not held conceptually but perceptually. And

what we're doing is bringing a perceived entity under a concept, we're identifying Fido, who we are looking at, by subsuming him under the concept "dog." Now, how can a child do this? On what basis can a child identify Fido as a dog? What is his justification for doing it? Well, it's simple. What does "dog" mean to the child? It means all those entities that fall within a certain range along a certain CCD. All the entities with a certain range of shapes, say, or a certain range of motion, which range of shape and motion it can distinguish perceptually. It had to be able to do that form the concept "dog" in the first place. All it needs to do to identify Fido as a dog is to notice that Fido's shape falls within the range of dog-shapes that he's already established in forming the concept "dog." And that's it. That's all there is to justifying a first-level proposition.

There's a big issue in contemporary philosophy: How could it possibly be the case that perceptual knowledge is justified on the basis of perception, which perception is not itself propositionally structured? This is a major issue for philosophers, but it's only a major issue because they don't have a theory of concepts that tells you how you can get to concepts from perception. So, really, once you have a theory of concepts, there's nothing to this, I think. It's just a straightforward application. All you have to do is notice that the units of a first-level concept have a characteristic which falls within the range of the distinguishing characteristic along the CCD. All you have to notice is that a dog's shape falls within the dog range. The structure of what one does when one does this is basically a deduction, but the premises are held implicitly. When the child forms the concept dog, he's saying, "I shall identify as 'dog' anything that falls within this range of shapes." Then he notices, "Fido there falls within this range of shapes. Therefore, Fido is a dog—I shall identify Fido as a dog." And that's all there is to it. Of course, he doesn't actually go through the deductive steps, but you can spell it out that way. That's the kind of reasoning that's taking place.

3.1.3. The unit perspective at the first-level of conceptualization

Now, let's talk a little bit about what the child has as a result of now having identified Fido as a dog or John as a man. The concept gives him an automatized unit perspective such that all the knowledge from each unit can be brought to bear on all the others. Now that he's identified Fido as a dog, all his knowledge about Lassie and Rex and Old Yeller is available to be applied to Fido. Now, what knowledge does he have about Lassie and Rex and Old Yeller (or whatever dogs he's seen in the past)? Since we're talking now about first-level concepts, when a child's first forming them, the knowledge he has is going to be perceptual knowledge. So, when you first form and start using first-level concepts, what they enable you to do is bring perceptual knowledge about past units to bear on the current units, the ones that you now have before you. Later, one is going to conceptualize what one knows about the units of one's concepts, and then one is going to have advanced conceptual knowledge to bring to bear on each of the units. So, now when we identify someone as a man, we have all the conceptual knowledge that constitutes the humanities (or which used to) to bring to bear on him. When you just identify something as a dog and you're two years old, you're not in that position.

But, just to give you a sense of what it means to hold information in a concept in a perceptual form and to be able to use it, let me give you an example. Suppose I told you that I once knew a horse-faced girl who moved like a cat. (I didn't.) Those are both expressions: "horse-faced," and "moved like a cat." You know what that means. "Horse-faced" is not a concept. It draws on your knowledge of horse and tells you, you know, "think of a face that looks kind of like a horse's." And what do you? You kind of stretch out the face and you make a long, flat nose, and that's what it means to be "horse-faced." We hold tons of knowledge in our concepts still in a basically perceptual form. That's why metaphors work; that's why there's such a thing as poetry, and so forth. And it has to be this way because, in fact, perception is an almost inexhaustible font of information. As you start conceptualizing, there's still a lot more to get out of it, and that's important for understanding the

progress of science, but that's a point later.

So our first-level concepts are file-folders, we can say. But they don't contain sentences. What they contain is is perceptual information about the concretes that fall under them—interrelated perceivings. And to identify something under a first-level concept (but also really under any concept) is to grasp what the thing is by relating it to other things. This begins right from the perceptual level, when we begin relating perceived entities to one another as wholes, and then integrating our knowledge of them.

3.1.4. Induction as beginning with first-level generalizations

We've already seen the role that concepts play in induction, but I want to bring that out; we've seen it implicitly. Notice that if we've already said that a concept is future-looking, and to form a concept is to institute a policy of applying what one knows from the study of *each* instance to the study of each other instance, to regard the instances as interchangeable, at least within a certain context, within a certain, you know, varying in degree. And this policy applies to information yet to be discovered, as well as to the information one already has about the units when one forms the concept. So to form a concept, the very formation of a concept, demands that one induce.

We've seen that when you form a concept, you're putting yourself on a policy of interrelating things, of omitting their measurements, of omitting measurements even that you haven't yet discovered, so that when you form the concept "man" you already know about, say, men's shape, and you're omitting that, but you're on the policy, "whenever I discover anything else about man, I'll omit the measurements of that, too, and hold it under the concept 'man.'" What that means is that whenever you discover something about any particular man, you do not view it as an isolated fact about that particular man. Rather, you omit the measurements of what you discovered about the particular man, and apply it to all men. Concepts demand induction—that's part of what it is to form one. And I think that this is what Dr. Peikoff means in his lectures on induction by calling a concept a "green light to induction." A concept is a policy of regarding each unit as one of unit a group of indefinitely many, and of applying what one learns about each to all of them. For example, if you have the concept "ball," you regard each ball as *a* ball, and everything you learn about the ball becomes ball-knowledge, which is to be applied to all other balls.

Of course, we do this with cognizance of the fact that the other units are different from the observed unit, but that they fall within the same range as it. An object *exactly* like this ball—the same shape, the same mass, the same size, the same material, the same everything—would behave in *exactly* the same way: roll in the same direction, etc., if it were in exactly the same circumstances—same force applied to the same point, etc. But something that's in the *range* of this ball, in the range of the circumstances that the ball is in, would behave in the range of ways that this ball behaves. Ball-range objects and push-range circumstances will exhibit rolling-range motion. The problem of induction, I think, is really the problem of finding the right level of generality—coordinating the level of generality for your subject and predicate. You have the concept "ball," and now the question is, "At what level should I identify the action that I see the ball performing?" If you said, "Balls roll at one mile per hour," that would be too narrow a generalization, and it wouldn't be true. But you can go narrower than "Balls move." In a lot of cases, it's very difficult to match up the ranges. You need to perform experiments, use Mill's methods, etc., and Dr. Peikoff talks about how that works in his lectures on induction. But in other cases, one has enough information already in the perception to interrelate the ranges, and I think that's the case with "balls roll." (We can talk more about this in the questions if people want, I just want to give a pattern for what happens here.) And so there are first-level generalizations—generalizations like "balls roll" that you can form directly from perception. That is, if you have the concept "ball," you can notice its action, you can interrelate it to the other actions of other

balls, form the concept "roll" (we'll talk about how you form action concepts in a few minutes), and then you just have validated the generalization and that's all there is to it.

Notice that the information we need—the information about causal relationships between ranges of characteristics—is the very information that we've already we need to form concepts in the first place. Concept-formation is itself an inductive act; it presupposes a host of implicit inductions, and it underwrites and demands future inductions, explicit and implicit. Similarly, as we've seen, the application of a concept—the making of a judgment—is a basically deductive act, and it supports and demands further deduction—that is, the deduction of all one's general knowledge about the kind denoted by the concept to the particular. So if I apply the concept "human" to Betsy, and say that she's a woman, I now have sanctioned and demanded that I apply all my other human knowledge to Betsy, so I know she's rational, she's got a heart, and, you know, so forth—everything else I know about human beings—she needs to be free in order to survive, and so forth.

3.1.5 The value of first-level concepts

Now, I want to say something about the value of first-level concepts. They're only possible because we know quite a lot already at the perceptual level. If we didn't have quite a lot of knowledge available at the perceptual level—knowledge of similarities and knowledge of causal relations—we would not be able to form first-level concepts. If we didn't both have this knowledge and be able to use it, in fairly complex ways: to compare, associate, project, etc., it would not be possible to form first-level concepts. And, in fact, much of—though not all—much of what we do with first-level concepts can be done without them. Animals can learn quite a lot, and I think they can even project the very near future, to a reasonable extent, based on associations, similarities—what we called visualization. An animal can't form the general proposition "balls roll," but animals have no trouble playing with balls—even new balls that they haven't perceived before—that is, if you give your cat a new ball.

What, then, do first-level concepts add? They add the consolidation of knowledge that can be perceptually acquired. First-level concepts' primary value is the consolidation of perceptually available knowledge. The formation of the concept is the formation of a new mental entity, a single mental unit, a new file-folder, a new subject of knowledge. Remember that we said propositions have a subject-predicate form. When you form a concept, you have new subject of knowledge: a new thing that you can know *about*. Before forming the concept, all of your perceptual knowledge about dogs is held in the form of various memories of particular dogs, more or less associated to one another, and if you're going to use that knowledge, you have to woozily associate all these things—hold one or two in mind at once, and project from those. What the concept "dog" does is it consolidates that into knowledge about a single subject: dogs—a single mental entity. And this enables you to use the knowledge much more unit-economically; it enables you to hold knowledge in the form of generalizations about integrated groups, as opposed to in the form of more or less associated memories about particulars. This, in turn, enables us to perform more complex acts of interrelation than animals can perform. A child and a dog can both see the similarities between two men. But because the child consolidates this knowledge with the concept "man," he is able to go on to observe further similarities that the dog cannot notice. For example, the similarity between men and dogs, between both of them and plants, and so on.

3.2. Higher-Level Conceptualization

And this brings us to the topic of higher-level conceptualization—that is, the formation and use of concepts which can only be formed and used on the basis of prior concepts. Quoting now from the beginning of "Abstraction from Abstraction," that is, one of the chapters in *Introduction to Objectivist*

Epistemology,

> Starting from the base of conceptual development—from the concepts that identify perceptual concretes—the process of cognition moves in two interacting directions: toward more extensive and more intensive knowledge, toward wider integrations and more precise differentiations. Following the process and in accordance with cognitive evidence, earlier-formed concepts are integrated into wider ones or subdivided into narrower ones.

Notice that all concepts are based on cognitive evidence and formed by a certain process. The cognitive evidence is: knowledge of similarities, differences, and causal relations. First-level concepts are those concepts the evidence for which can be held in perceptual form. Higher-level concepts are those the cognitive evidence for which must be held in conceptual form, thus these higher-level concepts logically presuppose the concepts necessary to grasp the evidence on which they are based. And we'll talk more about that presupposition relation in a few minutes.

3.2.1. The formation of higher-level concepts

Now, in this quote, she mentions two ways of forming higher-level concepts based on earlier ones—that is, narrowings and widenings. In fact, there are more ways than these two, and she acknowledged that... that this was not meant to be exhaustive. I'm going to talk about one or two other ways, and I think that there are probably more than I'm going to bring up as well.

3.2.1.1. Widenings and narrowings

But let's start with widenings and narrowings: wider integrations and more precise differentiations. First, what do we mean by these? What are examples? Well, an example of a widening would be "dog," "bird" and "fish" being integrated to form the concept "animal," or "animal" and "plant" being integrated to form the concept "organism." Each concept—that is, concept like "dog" and "bird"—is treated as a concrete, as a single unit, and it's integrated to form a wider concept, say "animal," in more or less the same way as the concrete dogs were initially integrated to form the concept "dog." In a narrowing, or what she calls a subdivision or more precise differentiation, what is an example of this? Well, "parrot" is a subdivision of the concept "bird," and "Senegal parrot" is a subdivision of the concept "parrot." In these cases, the distinguishing characteristic of the original, wider concept, say "bird," becomes a CCD, an axis, a longer line, along which one can distinguish subtypes. We'll talk more about this process in a moment or two.

3.3.1.1 The hierarchies of knowledge and of generality

What I want to do first is make a few points that are just raised by the fact that we can go wider and we can go narrower. There are some things that need to be distinguished. The fact that we can proceed from earlier concepts and base later concepts on them gives rise to a logical order to concepts: that some concepts presuppose other concepts. It gives rise to a hierarchy of knowledge in which some later knowledge depends on some earlier knowledge. The fact that we can go both wider and narrower gives rise to a second sort of hierarchy--a hierarchy in order from widest to narrowest. And these two things need to be distinguished. There are two hierarchies that come about as a result of this process. Now, in general, the concept "hierarchy" denotes any system with a certain structure. What structure? Well, here's what the Oxford English Dictionary has to say. It says: A hierarchy is, "a body of persons

or things ranked in grades, orders, or classes, one above another." I talk more about this and about different sorts of hierarchies in my hierarchy class, particularly the first lecture. Okay, we want to distinguish between the hierarchy of knowledge, which was also called the "logical hierarchy" or "logical order," and the hierarchy of generality, which we could also call a "taxonomy." And I have, in Exhibit C on the handout which I gave out yesterday, two... you know, I show both of these hierarchies with select concepts to distinguish them. If we go in terms of generality, "organism" is wider than "animal," which is wider than "mammal," which is wider than "dog," which is wider than "poodle." That's the order going from general to specific. But if we think in terms of what concepts depend on one another, "dog" is probably presupposed by "animal," so "dog" probably comes first with "animal" second. "Organism" certainly comes after "animal." "Mammal" probably comes after "organism," and certainly comes after "animal." So it's not the same. Sometimes a wider concept will be before a narrower one and then after a wider one, and so forth. You can look at the chart for evidence on that.

Now notice that she talks about the two *interacting* directions—wider and narrower—from the concepts that you can form directly from perception. How do they interact? Well, sometimes we need to widen in order to narrow, or narrow in order to widen. For example, you need to widen from plant and animal to organism before you can discover the existence of bacteria, which then allows you to form the narrower concept "bacteria," and then still narrower concepts for subtypes of bacteria. Or, you need to narrow from the concept "man," meaning "human being," to get the concept "man," meaning man as opposed to woman and the concept "woman," and you need to narrow from "deer" to get the concept "doe" before you can widen from a bunch of these concepts to get the concepts "male" and "female," which are widenings from narrowings from concepts for particular animal species. Okay, so that's the basic pattern.

3.3.1.2. How widenings depend on prior concepts

What I want to talk about now is how widenings depend on earlier concepts, then how narrowings depend on earlier concepts, and then on some other types of concepts, besides widenings and narrowings, that also depend on earlier concepts. Why do widenings depend on earlier concepts? Why can't we form "organism" straight off the bat from perception? Well, the similarities between a particular terrier and a particular redwood tree are just too remote to observe perceptually. They don't *look* alike. They do engage in a similar sort of action—self-sustaining, self-generated action—but not perceptibly so. It's only through the concepts "dog," "animal," and "tree," "bush," etc., that we're able to grasp enough about each of these type of existents in order to see the relationship between them. Think of it, uh... when I gave the example of the two fractions that were not directly commensurable because they didn't have a common denominator, but you could reduce them to 5/9 and 6/9. Now, you can't directly relate those fractions to one another, but there are operations you can perform on them, say dividing the numbers in the numerator and the divisor by whatever they were multiplied by to yield the reduced forms of the fractions. Once you performed that action, you can then interrelate them, and see that 6/9 is greater than 5/9.

Well, it's the same sort of thing here. You can't directly relate an oak tree to a terrier and say, "Oh, you know, they're both the same in that they both engage in self-sustaining, self-generated action; if they don't get food, they'll die; and they can reproduce." You can't *see* that. But, once you've integrated the oak tree with all other trees... the redwood, whatever it is, and you have the concept "tree," and that's a type of plant, and the dog's an animal and you can then see the relationship between animal in general and plant in general, and only through doing that can you see the relationships between the particulars. So, for the concept "organism," you need first to have concepts for particular types of organisms. To form the concept "animal," you need to first have concepts for particular types of animals. Now, you don't need to have concepts for certain ones—that is, you don't need "dog" to

form “animal,” or “cat” to form “animal,” or “mouse” or “bird” or whatever, but you need *some* concepts for *some* animals before you can form “animal.” Now, in fact, animal is sort of a tricky case, because it might be that you can form, you know, “four-legged-crawling-thing” first. The details, where it’s so close to the perceptual level, aren’t really that important. What’s significant and what’s clear is that you can’t straight-off grasp the relationship between some lichen and a seagull.

3.3.1.3 How narrowings depend on prior concepts

Let’s talk now about how narrowings depend on prior concepts. Individual entities, remember, are differentiated from one another, and from the background, in perception. So, in most circumstances, we have no difficulty telling individual entities apart from one another. There’s no issue, for example, in telling a parrot apart from a hummingbird—that is, a particular parrot that might be fluttering around this room from a particular hummingbird—you wouldn’t confuse them with one another. There’s no issue, in forming a concept, of telling them apart. The issue with regard to forming the concept “parrot,” is how you can see the two parrots as similar to one another in contrast to the hummingbird. Now suppose you were just looking at the two parrots in isolation. Let’s concretize this. My mother has a parrot; it’s a Senegal parrot. It’s green, about 7 inches long long, it’s inarticulate (my mother thinks it can talk, but it barely does), small and green. There are also Amazon Gray parrots, which are about twice as big, they’re gray, and they talk much better and much louder. If all you knew was these two birds—if, so to speak your mind was just an awareness of the Senegal parrot and an awareness of the Amazon Gray—you would see them as different, rather than as similar. To see them as similar, you need a foil—for example, Lassie, or a tree. If you bring Lassie or the tree onto the stage, suddenly the parrots look a whole lot alike. Now, imagine, in this context, that one brings in a hummingbird. The hummingbird clearly goes with the parrots as opposed to the other things. Even if you form a concept after you’ve only seen the two parrots and Lassie, and say you call that concept “parrot,” or “bird” or whatever word you use, as soon as you see the hummingbird, it would *have* to go with those other things. It would have to fall under the concept, because your context is: the way that the two parrots are similar to one another as opposed to the tree, and the hummingbird falls in the range that sets up by contrasting those two parrots with one another, as opposed to the tree.

Now, this is not a matter of what one happens to see when—that perhaps one happened to grow up in a parrot shop, or something, and saw a lot of parrots before one saw any hummingbirds, or whatever it is. (It would be a nightmare growing up in a parrot shop; I don’t like my mother’s parrot.) The context in which we form a concept is not the things that happen to be right in front of us at the moment we’re forming it. The context that we use when we compare things includes all of our perceptual memories, etc. Thinking of the units and the foil as being up on a stage is a helpful heuristic in some ways, but it’s misleading. In order to form the concept “parrot” as opposed to forming the concept “bird,” even if what we’re looking at is two parrots and a foil, we need a context in which to differentiate parrots as such from other existents, and this context can only be provided by the concept “bird.” Once you have “bird,” you can say, “okay, now we’re talking about birds.” Within that range (we ignore everything else, because we can focus just on the birds), and now we can say, “Look, aren’t these two different parrots like one another as opposed to the hummingbird?” If you didn’t have the concept bird to delimit the context in which you were talking, dog-thoughts and tree-thoughts would keep coming in, and you would not be able to hold that similarity. That’s why narrowings presuppose earlier concepts. That is, that’s why it’s not the case that you immediately form the narrowest possible concept for any group of existents, and then just widen from there.

So, we’ve seen how some concepts presuppose narrower concepts from which they’re widened, and some concepts presuppose wider concepts from which they’re narrowed. In both cases, sometimes other concepts are also presupposed—that is, in both of these sorts of cases, concepts for the

distinguishing characteristics, but not always, and I'm not going to talk about that, *per se*, now, I just want to point out sometimes that's true. You can ask me about it in the question period if you're interested.

3.2.1.2. Concepts of characteristics

I want to turn now to another type of concept that requires prior concepts to form, other than a widening or a narrowing, and that's what I'm going to call a concept of a characteristic. A characteristic is, in effect, anything, that's not an entity, that exists. Everything that exists is either an entity or else it's an action of an entity, an attribute of an entity, the relationship in which an entity stands, a quantity that some entity has, or so forth, and there are lists—Aristotle made a list of how many things like this there are, and one can quibble, but I don't want to quibble, I want to focus on the distinction between entities and other things. All of those other things exist *of* an entity: action of an entity, attribute of an entity. And "characteristic" is the widest term I'm using to subsume all of these. So it includes attributes like "red," and "long," actions like "walking," relations like "underneath" or "being the father of," quantities, etc. I want to note, before going any further, that characteristics can have characteristics themselves. Adverbs denote qualities of actions or attributes of actions. Comparatives like, "bigger" or "more" denote relations between qualities or quantities. There might even be actions of actions—acceleration, perhaps? We don't need to settle all of that, but characteristics can, at least in some cases, have characteristics.

3.3.2.1. How concepts of characteristics presuppose earlier concepts

Now, the concept of a characteristic presupposes concepts for the existents that it characterizes. A concept for a characteristic presupposes concepts for characterized existents. One doesn't necessarily need a concept for every existent that it can characterize. In order to form the concept "blue," one doesn't need a word for every entity that can be blue, but one needs at least *some* concepts for *some* blue entities before one can form the concept "blue." One could not have adjectives without concepts of entities. One could not have verbs without having nouns denoting things capable of performing the verbs. One could not have adverbs without concepts for the verbs to which they apply—that is, verbs which are concepts for the actions to which the adverbs apply, etc.

Why are these concepts for characterized things presupposed by concepts of characteristics? Well, entities are what is isolated for us in perception. We're aware, perceptually, of the characteristics of the entities, but only as part of the whole entity, not in isolation. We can focus on them selectively to some extent, but we cannot fully isolate them and form concepts of them without prior concepts of the entities. Why not? Well, I think for two reasons: similarities and differences between the entities as wholes would always trump similarities and differences between the characteristics. Compare this to the case of the difference between all birds and horses trumping the difference between parrots and hummingbirds in the case of narrowings that we just talked about. But also, I think another reason why concepts for the characterized thing are presupposed by concepts of characteristics is that part of what one has to omit in order to form characteristic concepts is the things characterized. Those things have to be omitted as wholes, not as bunches of individual attributes, and, I think, in order to perform that act of omitting the entity in forming the concept "blue," you need a concept for the entity. So that's another reason why I think concepts of characteristics presuppose concepts of the things characterized. Ayn Rand doesn't write about this, that is, she doesn't write about why this is, but she does, in several places, make points like, "verbs presuppose nouns." I discuss this somewhat, and in some detail, actually, in my article on *Anthem* and why the concept "I" is needed, if people are interested in that, and I have references to where she discusses it.

[Betsy Speicher is apparently waving her hand about frantically]

GS: Um... I'm norm-... not trying to take many questions as we go, but uh, I'll do one now 'cause Betsy seems anxious.

BS [clearly audible]: Okay. Um, uh Ayn Rand describes a first-level concept, as held by an infant, as something...a concept of "man" as something that moves and makes noises.

GS: Mm-hmm

BS: Okay, that... that would assume that uh, the child obviously doesn't have a concept of the action or of the noisiness--the characteristics--, but it's like he perceives the characteristics to form the concept of...

GS [overlapping with her next words]: Yes, that's right

BS: of a kind of entity.

GS: I think in forming concepts, in fact she *says*, "in forming concepts of entities, you're aware of their characteristics," and it's their characteristics that you're differentiating them by. What you are *not* aware of is the characteristics as totally isolated things, about which you can think. They can't be subjects. I really do want to move on, though, can we hold this for later?...okay

So we're talking about concepts of characteristics now, and how they presuppose concepts of what they characterize. That is, I'm saying, "of what they characterize" because in the primary case it's entities—concepts of attributes presuppose concepts of entities—but the point applies more widely: adverbs presuppose verbs, for example. You know, concepts for characteristics of actions will presuppose the concepts for the actions, and so, in general a concept for a characteristic will presuppose concepts for the existents characterized.

3.3.2.2. Concepts of consciousness

Concepts of consciousness are an example—a special and important example—of concepts of characteristics. Conscious states are actions, or states, of a subject. And consciousness the faculty is an attribute of the subject. Therefore you cannot form concepts of consciousness—concepts for particular conscious actions like "see," "think," etc.—except on the basis of an awareness of the entities that can perform these actions, and some existents on which these actions can be taken. So, for example, to form the concept "see," you have to have the idea of *someone* that can see *something*. I see this cup. Earlier I was seeing this desk. Betsy can see me. Paul can see me and Betsy. There's seeing that all of us are doing. Some, but any, person can see some, but any, thing. And that's the concept "see." So that's why you need prior concepts to form concepts of consciousness, or that's, anyway, part of the reason. Concepts of consciousness falls of under this wider category: concepts of characteristics.

3.2.1.3. Concepts for conceptually discovered existents

Another category of concepts, besides concepts of characteristics, besides widenings, besides narrowings, that presupposes earlier concepts—this one, I think, is easer to get one's head around: concepts for conceptually discovered existents. For example, the concepts "electron," "microbe" and "photosynthesis." You don't know that the units of these concepts even exist except on the basis of knowledge that you got using concepts. That is, you can see blue when you see a blue thing, it's just that you need concepts of the entities before you can abstract it, but you don't have *any* form of awareness that there's such a thing as an electron, a microbe, or such a process as photosynthesis by just looking around at the world. Only once you've formed the concept "plant" and you've done a whole lot of chemical reactions, which require a whole lot of other processes, do you have any

knowledge of photosynthesis at all.

Okay, so we've now covered several categories of higher-level concepts: widenings, narrowings, concepts of characteristics—including concepts of consciousness—, concepts for conceptually discovered existents. We haven't said much about this last, but it doesn't seem to me that there's very much to say, I think it... it's clear.

3.2.1.4. Cross-classifications

One last issue about abstractions from abstractions that I want to discuss here. If we look on our handout under Exhibit C, on the right, we have select concepts ranked according to the hierarchy of generality. And so the most general concept in the right-most chain is "artifact"—that's, you know, man-made thing—under which we have the concept "furniture," because a "furniture" is a type of man-made thing. A "table" is a type of "furniture," and a "coffee table" is a type of "table." Simple enough: we go from wider to narrower. But what do you do about, say, "patio furniture," and "patio table"? Is "patio furniture" wider or narrower than "table"? Well, it's not like all tables are patio furniture, or all patio furniture is table, you can't rank them relative to one another. There's something else going on here; this doesn't belong at any particular place in the list from general to narrow that goes artifact, furniture, table, coffee table. It's something else, it doesn't belong in this chain. There's, rather, an orthogonal classification, a cross-classification that's going on here. After you get to furniture, there are different ways you can break furniture down. You can break it down into different pieces of furniture serving different purposes—you know, tables, chairs, etc.—or you can break it down into furniture in different locations.

Notice, if you think about Ayn Rand's definition of furniture—which I can't remember the full, exact wording for, but it specifies that furniture is in a human habitation, and that's it's used to support smaller objects or contain them—well, one way of narrowing further narrows on "what is its particular function, what smaller object is it used to support or contain?" and that gives you "table," "bed," "cabinet," etc., and another one is, well, "what sort of human habitation? A living room? A patio? And you get "patio furniture." In fact, this is what Ayn Rand calls a "qualified instance" rather than a normal subdivision, but I don't think that matters for this point. If people want to know about the difference between a qualified instance and a normal subdivision, you can ask during the question period—It's really immaterial for what we're doing now. On my little chart of the generality, I kind of took care of the ambiguity of "patio furniture" and "patio table" by making these little dotted lines instead of normal lines.

Well, what's happening here? if we think of the first-level concept, which would be "table," we can think of the relationship between "table" and "furniture," and between "table" and "coffee table," and between "furniture" and "artifact" are essentially the same type of relationship. It's the relationship that we capture by saying that the narrower one is a species of the wider one, and there are going to be a number of other species which are mutually exclusive with the species in questions—that is, something can either be a table or a chair, but it can't be both—and eventually you can get a jointly exhaustive list of them, probably, although that doesn't matter (or you might be able to—In a lot of cases you can't, but that's of no consequence). There are cases where there's just a different sort of relationship, the sort of relationship that's brought by bringing in a different sort of characteristic. And I don't know exactly what to say about this except that I think that the cross-characteristic just does not narrow the thing that it's being narrowed from in the same way. In particular, it doesn't narrow the thing being narrowed down *as a whole*; it doesn't isolate and integrate the units as a whole in the way, say, the concept "man" integrates men as a whole, or "table" integrates tables as a whole. It's more like what happens when you isolate an individual characteristic: you group the things together on the basis of the characteristic, thinking that a lot of importance will follow from this grouping by characteristic,

but not enough, and not in the same way (I'm not sure how clearly I can specify this) that is the case with a normal concept. And so you think of this as a kind of special case or a special way of classifying within the context of another, primary way of classifying. It's a derivative way of classifying. I mention this because I think it's important for certain topics in theory of induction and theory of definitions to distinguish these concepts from other sorts of narrowings and other things in a hierarchy of generality. People can ask about some of the issues that the distinction raises, if you want, in the question period, but I think that this is a point that's worth people thinking further about, so I wanted to raise it.

I'll just make one more point in this connection before moving on. In Aristotelian philosophy, there's the concept of an *infima species*; it means narrowest species. (Aristotle's literal words are "uncuttable form.") You can't cut it up further. I think this is a valid idea. That is, I think there are concepts that cannot be subdivided except as cross-classifications. For example: "man." I don't think you can narrow "man" to get anything narrower than man in the same way that you can narrow "table" to get "end table" or "desk." The thing you get is different, because the relationship between man's distinguishing characteristic and his other characteristics is such that the way in which that causes the derivative characteristics breaks down if you get too narrow. If you think about the essential trait of an animal, which is that it can perceive and locomote, and you think if you narrow down how it perceives and locomotes, a lot of things change with that and you get a kind of whole. A animal that perceives and locomotes differently from another animal is really a whole different sort of animal. That is, something that perceives by echolocation and flies around is a bat. If you change the mode of perception to, you know, vision and walking, you've got, you know, a dog or something; you have a totally different kind of animal because a lot changes with that. But if you narrow down man's essential characteristic—his reason—to, for example, "people who reason really well" and "people who don't reason really well," a lot does change: you get, in effect, the concept[s] "genius" and "idiot," and so you get "Aristotle, Ayn Rand, and Newton" as opposed to "Forrest Gump" or something. And a whole lot is different, and we do need a concept "genius" as opposed to "idiot," but a lot *doesn't* change that *does* change between, say, "bat" and "dog." For example, the genius' anatomy isn't very different from the idiot's. So, I don't have, really, that much to say about this, except that I think it's a topic worth some further thought, and I wanted to raise it for people thinking about this.

3.2.2. Dependence of Higher-Level Judgments on Lower-Level ones.

Okay, a few more issues before we get to the end for today. Higher-level propositions depend on lower-level propositions as higher-level concepts depend on lower-level concepts. The identification of a given thing as a unit of a given concept requires that one use the earlier concepts which underwrite one's subject concept and predicate concept. anytime you're using the concept "animal," that presupposes all the concepts that you need to form the concept "animal," and anytime you use the concept "organism," it presupposes all the concepts needed to form the concept "organism." For example, the concepts "bird" and "animal" are both at work in identifying a given bird as an organism because to identify the bird as an organism is to interrelate it to, for example, fungus and plants. And this can only be done through the intermediate identification of the bird as a type of animal—of *this* bird as a bird, bird as a type of animal, animal as a type of organism. You can't directly interrelate it to fungus. The identification of a bird, as being of the same kind as fungus, of being like fungus in that they're both organisms, proceeds through what logicians call "middle terms." That is, it proceeds through the terms "bird" and "animal". You can't think of the bird as a type of organism without first thinking of it as a bird and a type of animal. There is, in effect, a type of deduction involved: this thing is a bird, bird is an animal, animal is an organism, therefore this thing is an organism. Again, you don't go through all these steps—you don't truck through them, just as you don't go through the relevant steps like that in initially applying the concept "bird" to something just by noticing it falls within the

bird-range. But these steps, which you can lay out, identify a pattern of what goes on in your reasoning, and they're an abstract summation of what it is to identify something under a concept.

The same is true with identifying a particular dog as a poodle—that is, with using a narrowing of the concept "dog." One can only identify it as a poodle in a context in which one is already identifying it as a dog because, except within the context provided by that concept, one would not be able to see a given poodle's relationship to other poodles, *per se*—as opposed to it's relationship to other dogs. We discussed the reasons for this when we talked about why one wouldn't be able to form the concept "parrot" just by looking at two parrots and a horse. You can only relate the two parrots to one another and *not* also to the hummingbird in a context in which you've distinguished birds collectively from non-bird animals. Just a qualification on this: of course you can discover a new kind of animal that you've never seen before and grasp immediately that it's an animal before forming a concept for the type of animal that it is. Say you've never seen a sloth, and you go to the zoo, and you see a sloth, you can say, "Oh, that's an animal. I wonder what type of animal it is." But notice that when you do this, you're really not immediately grouping it under "animal" and thinking of this thing as a sort of *sui generus* animal, as *directly* a unit of "animal." Rather, what you do is you think of it as something that belongs under some other concept which you don't have yet, which *then* belongs under animal. You think of it as "some *kind* of animal. What kind of animal is it?" You don't think of it just as "animal" straight off.

So that's the point that higher-level propositions depend on lower ones. And it really is just the same point that higher-level concepts depend on lower ones, because a higher-level proposition is a use of a higher-level concept. The concept is an awareness of a group of things in relation to one another. The judgment (of which the proposition is the form of the content) is the relating of a new thing into this group. It's the bringing of a new thing into the group of interrelated things, and it requires everything that the formation of the concept requires. The formation of a concept is not like making a brick which you then build a house with, it's like forming a cell which is part of a living thing. The concept itself is an ongoing activity or process which has become institutionalized, and use of it requires all the things that forming of it requires—that is, the same kind of actions are taking place.

3.2.3 Logical order distinguished from chronological order

Now, we've been talking a lot about presupposing—what presupposes what—and I want to make a distinction here that Ayn Rand insists on in ITOE, and which I think is very important: the distinction between logical and chronological order. Chronological order is order in time: what happened before what else. Logical order is order of dependence. What's the difference? Well, I learned the concept "cat" before I learned the concept "dog" because there happened to be cats in my house, growing up, and no dogs. That's a chronological order—It's what one I learned first. I also learned the concept "cat" before I learned the concept "bacteria." I couldn't have learned the concept "bacteria" before I learned "cat" (*or* some other animal concepts at that range of generality, e.g., "dog.") In some cases, the concepts *depend* on the concepts that came before them in time. In other cases, one concept precedes another in time for some accidental reason—It didn't have to. Thus Dr. Peikoff describes logical order as "the necessary order of learning", as opposed to the order in which one happened to learn.

Logical priority implies chronological priority, but the reverse is not the case. This said, there are cases in which you can *kinda* form a logically later concept before you have a logically prior one. Kinda, sorta—you're not clear on it; it's confused. When exactly you can do this is not such an important question . As epistemologists we're not interested in people's chronological development—what they learn when—or even what they *can* learn when. That's an interesting topic. It's useful if you, say, want to be a teacher and design a curriculum, but that's not what epistemology is about. The whole

reason we want to know about the logical order is so that we can test and validate our concepts, as going to discuss tomorrow. Whatever might have happened when you were a kid doesn't matter for that purpose. The way you tell what concepts presuppose other concepts is not by asking or studying the order in which children learn them. Now, it might be useful as a heuristic or clue in some cases, but it's not the evidence, it's not the proof of the order. The proof is now, the proof is: Right now do you need to think about something as an animal in order to understand how it relates to other organisms? Can you, right now, straight off, perceive how the parrot that you're holding in your hand relates to a piece of lichen or relates to the cold that you're experiencing? Can you see it? No. Well, then, "organism" is not a first-level concept. What do you need to get your head around, the relationship between this parrot, this piece of lichen, and this cough? You need to see that the parrot belongs with the horse and the bird, and so forth—you need the concept "animal"; then, under the concept "animal" you can hold enough information to relate animals to trees, and then you can relate form the concept "organism", which enables you discover bacteria or viruses and see the relation to the cough. Okay. So that's all I want to say about logical order and chronological order. Logical order is what depends on what.

3.2.4. Induction

I want to move on now to saying a little bit about induction. We've spoken about it before in connection with first-level concepts and anticipated what we'd have to say later about it. Higher level concepts make it possible and mandatory to proceed from first-level generalizations like "balls roll" to higher-level generalizations. We've already seen how induction is a matter of coordinating the level of generality of a subject and a predicate, and I refer you to Dr. Peikoff's lectures on induction for a discussion of how this is applied in higher-level cases. I want to mention now, though, that induction is a really wonderful case of the way in which widening and narrowing interact. Remember we said that sometimes you have to narrow in order to widen, or widen in order to narrow; you need to form a more precise differentiation in order to be in a position to form a wider integration.

In Dr. Peikoff's lectures on induction—in physics in particular—he shows how mathematical concepts allow us to make more precise differentiations: to more precisely identify the actions we perceive—for example the motion of a ball when it rolls, or of a pendulum when it swings—so that instead of just saying, "this ball rolls," we can say, "this ball rolls with a certain force in response to impetus exerted on it in a certain amount," and, in general, come to up a formula for the way balls roll. The equations in physics are just complex judgments. It would take some philosophy of math to say exactly how that works, but I think the general idea is clear. The result of this is that precise, quantitative identifications of actions enable us to grasp relationships between seemingly disparate items and actions. For example, through a very long chain of inductions, all of which involve getting more and more precise mathematical formulas, we're able to interrelate the motions of falling apples and wandering planets. That is, we're able to have Newton and then Einstein, presumably, although I don't really know how any physics after Newton works (and I'm shaky even up to him). So I just want to sort of suggest that as a theme to have in mind when thinking about Dr. Peikoff's work on induction, and to show sort of how it relates to some of the other issues we've been talking about. Theory of Induction is a whole sort of complex topic in itself.

Dr. Binswanger?

HB: You said "precise qualitative," I think you meant "precise quantitative."

GS: Yes, preci-... sorry, did I say qualitative? Wh-...uh, the math gives you a way to quantify precisely the relationships between, say, uh, uh..... it allows you de-... to li-... delimit strictly the range of measurements involved.

Allan Gotthelf [faintly]: ...and say how this in-... this process involves a narrowing and a widening?

GS: Sure, the question is (and this'll be the last question before I go on a little, and then I'll stop again for questions), "How does this process described involve a narrowing and a widening?" Well, you first...uh, you perform your experiment, let's say, uh... even before you perform your experiment, but let's imagine you're... you have a pendulum experiment going, and there's a kid in the room, and he looks at it and he can form a first-level, uh, description of what's going on, "That thing's swinging." You go, you know, "that thing's called a 'pendulum'"--okay, he's got the concept "pendulum"—"Pendulums swing." Uh, what the scientist does is he, uh, forms a narrower description of that--not just "pendulums swing," but "pendulums of certain lengths swing in certain ways with certain mass bobs, and so forth." Um, so he, uh, doesn't just say, "pendulums swing," but, "pendulums of this sort swing this way." That's the narrowing. And the "this sort" and the "this way" are, you know... stand for the mathematical expressions that he uses to describe the pendulum and its motion.

Um, he then does, or o-... he, or other scientists, do similar things...uh, similar experiments and observations which involve quantifying um, uh... quantifying, and thereby narrowing, earlier observations. So, for example, you look up at night at the stars, and you see some fixed ones that kind of orbit together, and then some ones that wander, and you form a concept for that, "the wanderers," which is just what planet is—*planein* is "wander" in Greek an-...okay. So, now you have "planets wander," which is a low-level generalization --probably not quite first-level. Well, then somebody f-... narrows that down, you know, "Mars is wandering, it's erratic movement is actually this particular, narrower movement, which, uh, and ...that's... it wanders *thusly*," where the "thusly" stands for a kind of mathematical description you have of Mars' motion.

Now once you've done a lot of this, you-... you know--the steps is the history of physics--you start to be able to see relationships between the two "thusly"s, that is the "thusly" for the way the pendulum moves and the "thusly" for the way the planets wander. And now you have a wi-... you're able to go from that to a widening--not about how pendulums or planets move, or whatever the intermediate steps may have been--but about how massive bodies move, and that's Newton. Um, and there are all sorts of steps, you know,--I'm just sort of forecasting in broadest terms what the pattern is, and I think Dr. Peikoff, uh, ver-...you know, illuminatingly goes through how this is the way to understand various stages in the progression of science, um... culminating in Newton, um... although he discusses other ones as well. Um, so I think that's one of the... the most, uh, interesting lessons of that lecture series--at least it was to me.

3.3. The Nature of Conceptualization

I want to talk a little bit about the nature of conceptualization in general, abstracting from what we've been talking about, and about reason, before stopping for questions today. Consciousness is an active process, consisting in two fundamentals: differentiation and integration. This applies to all consciousness, not just conceptualization. Sensations, as we discussed two days ago, are primitive differentiations. Perception is the physiological process of integrating sensations to form an awareness of entities differentiated from one another and from their background. (Or else, the word "perception" can refer to the awareness we have through this physiological process.) Concepts are integrations of perceptions, as perception are integrations of sensations, but concepts are forms in which we are aware of *kinds* of things as differentiated from one another; perceptions are forms in which we are aware of *things* as differentiated from one another. Sensations are forms in which we're primitively aware of differences—yum or ick—except *we're* not primitively aware of differences in that form. (Maybe a

scallop is.) The move from sensations to percepts is done physiologically, at least in human beings, by the brain, so that we start with percepts. So that's one difference: the integration that takes place in conceptualization is a consciously performed integration, and it's a volitionally performed integration (the volitional part we're going to focus on tomorrow).

3.3.1 Integration by interrelation

But I want to make just a more general point about integration and interrelation that applies to the move from sensations to percepts and from percepts to concepts. Consciousness throughout, anywhere in this stage—If you're at the sensation level, if you're at the perceptual level, if you're at the conceptual level—is a process of identification. And we identify things by relating them to each other—basically, by differentiating them from one another. To know what something is is to know how it differs from other things, to grasp "this" versus "that." Put it another way: if you want to identify something, what are you going to identify it terms of? There's nothing but what is. So the only thing you can identify "this" in terms of is other things that are. You identify something by differentiating it from something else, and a sensation is just a differentiating of something from something else at this most primitive, scallop-like level: the differentiation of the yummy environment from the icky environment.

A perception is a differentiation of one entity from another entity, and this differentiation is achieved by the integration of the information that's already, in some form, in the sensation. There's some kind of processing or correlating of the sensory information of whatever comes in to our senses; the energy that stimulates our sense receptors is integrated, coordinated mechanically by the brain to yield a differentiated, discriminated awareness of entities. Everything we're doing after that is the same kind of act, at a very general level. What we're doing is: being aware of things that are different from one another, by using our awareness of their differences to grasp more about what each of them is. You grasp what each thing is by grasping how it differs from each other thing, and how each part of it differs from the other parts of it—each attribute, aspect, etc. So differentiation is really what's going on here. But there's just too much information to grasp all these differences sort of individually. One differentiates wholesale, one differentiates in an economical way, by integrating. What one does is one puts together the things that are not so different so that one can differentiate them collectively from the things which they're more different than. And then one applies what one learns to the individual ones by differentiating them again from one another.

3.3.1.1. Integrative interrelation as performed in a certain logical order and as operating at tiers of generality.

Identification in general is interrelation. It's differentiating things, integrating them on the basis of one's differentiations, which integrations allow for further differentiation. The purpose of the integration is to condense the data to allow us to differentiate further and to hold our differentiations better. And this has to be performed, this differentiating and integrating, in a certain order because of how much information we can hold at once. And it has to be performed at different levels of generality. So learning is a process of sort of stepwise integrating—of differentiating things, integrating them based on the differentiations one made, which allows further differentiations, which allows further integrations, in steps. Thus the logical order, the logical hierarchy, the hierarchy of knowledge. And this all has to be held in tiers of generality, thus the hierarchy of generality. That is, if I want to fully identify Paul—to identify him by differentiating and relating him to all the other existents—I have to think of him as a man, which is a type of animal, which is a type of organism, which is a type of entity, which is a type of existent, and then, further, I need to think of him as a particular type of man. Now I

say there's no subdivision of man in quite the way there is of animal, of animal to get man, but he's a doctor, he's, has brown hair, etc.; he's wearing blue, wears clothing rather than a primitive savage who runs around naked; there are different ways in which I can differentiate him.

So these two hierarchies are both significant. The hierarchy of generality allows us to hold the integration of all of our knowledge, hold how each relates not just to the things it's immediately differentiated from, but to the whole—that's what the hierarchy of generality is for. And we've already glimpsed the role it plays in induction because you want to see how high up the role of generality, or how low down it, the hierarchy of generality, you have to go to get your subject and your predicate. The logical hierarchy plays a different role. Having explicit knowledge of the logical order—and I'm going to talk more about the need to have explicit knowledge of it tomorrow—lets you see where the knowledge came from, how you got to it, and it's going to be, as we'll see when we talk about objectivity, crucial for reduction.

3.3.1.2. The role of axioms in integration

In this connection—and this is all I'm going to say about axioms, other than what I said in the question period—axioms are concepts that are all the way up on the hierarchy of generality. They're the most general concepts that allow us to integrate all of our knowledge. And they're at the very beginning of the logical hierarchy. That is, the information that's captured in axiomatic concepts is available directly in perception, though we don't form the concepts until we recognize the need to hold that information in a certain form in order to play a certain function in our cognition, and that function in our cognition is integrating the whole of our knowledge—standing, in effect, at the top of the hierarchy of generality.

3.3.2 AR's definition of "reason"

I just want to say a few words about reason, and then we'll stop for questions. Ayn Rand defines reason as the "faculty that identifies and integrates the material provided by man's senses." Reason is a faculty for processing perceptual data. We've already seen how concepts are based on perception, how they're integrations of perceptions. And this is really the basis of the point we discussed two days ago about man being born *tabula rasa*, about the senses being the only basic or ultimate source of information.

3.3.2.1 AR's view contrasted with rationalism and mysticism

We're now in a position to understand how conceptual knowledge is just a processing of perceptual knowledge, we can now see that it's arbitrarily to claim that there's any source of knowledge independent of perception, and moreover that it's nonsensical to claim it. What it is to know something is to interrelate perceptual data. So the idea of some knowledge that doesn't have its basis in perception is just incoherent, it couldn't be anything like our other knowledge; it could stand in no relation to anything else we know. So this is a refutation of rationalism, that is, the idea that there's some knowledge that's gained independent to, and is therefore somehow superior to, knowledge gained in perception.

3.3.2.2.AR's view contrasted with empiricism

But the point that reason is the faculty that identifies and *integrates* the material provided by man's senses shows that Ayn Rand is not an empiricist either. Empiricism, as that concept is used in

philosophy, and as I think it ought to be used, given its origins (if people want to know about that they can ask), refers basically to the view that all our knowledge *is* perceptual knowledge, not is *based* on it—that there's no act that we perform that gets us from perceptual knowledge to knowledge that is essentially different from perceptual knowledge, that knowing is, in effect, perceiving. Reason is the faculty of identifying and integrating the material provided by the senses, and those of you who have studied the history of modern philosophy, in particular, will know that there's a distinction between rationalists and empiricists; Objectivism is neither. I think it's more obvious how it's not a form or rationalism, but people sometimes think of it as a form of empiricism, and that's not true, it's just essentially different from both.

Q&A

Don [faintly]: [Could] you say something about the, uh, qualified instance for sub-classification that you mentioned?

GS: Yeah, um... Coffee... this is "coffee table" versus, um... suppose we had a, another name for it, um... Well, it- it's the... the diff-..... "Poodle" is a subdivision of "dog," um... "toy dog" is a cross-classification of "dog." Um... "coffee table" is a qualified instance of "table." Um... what's the difference? Well, a qualified instance is one where I think it is a narrowing of the concept, but... the... sort of, it's not important enough to set up a whole independent, uh, file folder. So you don't store the inf-... you don't sort of have a separate data bank about coffee tables (unless you're Kramer or somebody who's writing a coffee table book about coffee tables). I mean, you might, that is--the point is that you might in some contexts have a separate sort of knowledge bank about coffee tables, but not normally. You file it as, as knowledge about tables, uh, with some special stuff added. So it's, um... Under your knowledge about tables would be, "they're of different heights--if they're for serving coffee on, they're pretty low and they're by the couch in the living room," rather than, "table, subtypes: coffee tables, kitchen tables, dining room tables--see these other folders," uh, and then in the "coffee table" folder you have the knowledge about coffee tables. Now that's just an analogy, you know, it's using the file-folder analogy, but I, I think it's clarifying, um... uh it's in "The Cognitive Role of Concepts," uh, I think primarily... I think it's in that chapter that Ayn Rand discusses this a little, and so I, I refer you to that. And also somewhat in the chapter on definitions.

Um...I don't know, who else... go ahead.

?? [faint]: What do you think of the idea that the distinction between cross-classification and narrowing, widenings, is, um, that the conceptualization is based on non-essential characteristics for a specific purpose?

GS: Um...

??: ...versus essential characteristics for...

GS: I think... the question is, uh, is, "Is it the case, or what do I think of this proposal, and the proposal is that the difference between cross-classification and other widenings is that the differentia or the distinguishing characteristic in the cross-classification is a nonessential." Um, I don't quite like it; I don't think that's right, because essential is, uh, relative to our need, uh, that is to- to the cognitive need...um...and it requires that there... any concept requires there be some kind of fundamentality, where one thing explains a number of others, uh, and that's going to be true in cross-classification, that somebody's a doctor is going to explain a number of things about him, for example--"doctor" is a cross-classification as a narrowing of "man." Um, and it is essential, that is, it does serve the same function for us that, um, uh, that reason does in, in, in forming the concept "man." Uh, it has to do, I think, with the... whether all of the characteristics are narrowed. That is, when you take the concept "animal" and you narrow it to "man," right, the concept "animal" denotes a range, and there are ranges

along many many axes of measurement. When you narrow it to form the concept "man," almost all those ranges narrow accordingly. The shape changes to the man shape, the range or awareness changes to the man range of awareness, the type of movement changes to the man type of movement, the coloration changes to the man range of coloration, which is narrower than the animal range of coloration, and so on for practically any dimension you can think of along which you might measure men and animals. When you narrow "man" to "doctor," a whole lot goes along with it, but not everything. Nothing about the anatomy narrows, nothing about the shape narrows. What narrows? Well, what they think about, how they spend their days, how much money they have, how they get the money, and so forth--uh, how likely they are to be coerced by government officials, you know, all these sorts of things narrow, but not everything. And so I think...uh, I think that's the difference. But this is something that I'm sort of early in my thinking about, um, okay...

Uh, Dr. Binswanger?

HB [clearly audible]: I'd just like to say, I don't think there's any such thing as cross-classificational *widening*, it's always a narrowing

Unknown male: Yeah, I didn't, I misspoke.

HB [more faintly]: Oh, you said the narrowings...

GS: I'm not convinced that that's true. I'm not sure that it's *not* true--that is, is it the case, uh... you see, it's very hard to think about whether there are cross-classificational widenings. What might be an example of one is "male." If you think of "male" as a wide-...or "female" ('cause it's easier)--you think about "female" as a widening from "woman," "doe," "hen," etc., but you can also think of it as a narrowing from "animal" that's kind c-, uh... analogous to

HB [was speaking in background too faint to make out, at this point becomes much more audible]: ...classification...it is a widening

GS: Yeah... oh, I see, that's true, it's a wi-

HB: ...widening, it just summarizes. It can't be cross-classifi...

[4 second pause]

GS: Um... let's continue to talk about this after class. I have some examples of them, but they're kind of technical, and I'm not sure that's the right way to understand them. One, though, for example, is "belief," uh, which I think is a widening, of a sort, from knowledge. Anyway, it's a concept that requires knowledge as its base, and it is wider than knowledge, but it's formed in a kind of s-... un-... a way that's different from the way that we've been discussing with these other concepts, and, uh, I'm not quite sure how to... how exac-... I have thoughts about how it's formed, and other concepts that are formed like it, whether it ought to be understood as a cross-widening or something else, I- I've often called these concepts "non-generic widenings," um... but that just tells you what they're not, not what they are. Um... so anyway, there are, um... there are cases that you might describe as cross-classificational widenings. Um, I'm not sure that that's the best way to describe them. This goes to the more general point that there are a lot of interesting things to think about here. There are more types of concepts which require other concepts to form them than you might think at first.

Um, go ahead Gina.

?Gina [accented, soft but audible]: ...the difference of the relationship between the... the genus of a definition and the conceptual common denominator?

GS: The question is, "What's the relationship between the genus of a definition and the conceptual common denominator of a concept?" Uh, the genus of a definition denotes a group of things. The conceptual common denominator denotes an... a characteristic. The genus is the group of things which, uh, include all the units of the concept, and also other things which share a conceptual common denominator with them. So it's the group of things from which the units of the concept are

differentiated along some conceptual common denominator. So, for ex-... -xample, uh, well I think that's...[trailing off] that's sufficient for example.

Uh, Jean? And then David.

JB: I'd like to ask you to talk a little bit more about this idea that in order to form a concept like "parrot," you have to have the concept, "bird." You said that it sort of helps you delimit the subject so you don't start thinking about dogs and trees. Are you... are you saying... are- are you basically saying it gives you a mental set for that or a psycho... psycho-epistemologically it's getting warm in the right material of your subconscious, or [the rest is covered by your next words]

GS: Uh, Jean asked me, uh, the... what... w- to elaborate on the point that, uh, we need the concept "bird" to form the concept "parrot," uh, because it prevents us from, uh, automatically bringing in irrelevant things like dogs and horses. Uh, she then gave two sort of things I might say: one is that it "puts us in the right epistemological set"--yes, I agree with that. Uh, two that it "gets us warm...uh, in the right area, like you're getting warmer, you know, an- I think that's what... I don't think that's... oh, warms up the right databank?" Uh, I don't think that's really it. It probably does do that, but I don't think that's the essential. I think the essential is that it sets a domain; it says, "We clearly have differentiated bird from other animals, now it says 'birds': we're thinking about birds. So nothing that's not a bird is gonna come in, you know, to this new concept that we're dividing. We've distinguished birds from other things, now within 'birds,' um, (so don't bring up horses and stuff), uh, let's start looking at the similarities and differences between things." And that's what I think it does: it delimits. So it's not an issue of what things are near to the surface, what things are ready to be accessed, although I'm sure that's true, that it does make certain things ready to be accessed and not. I think of it more this way: it provides a criteria for what's in the domain that we're talking about, and what's not, rather than thinking of it in terms of psychological mechanics or what things it brings to mind or not, although surely it must do something with the mechanics of psychology, too.

We're running out of time, and I promised I'd do David, so let's do him and then, last[?]

David[faintly]: Uh, if first-level [now louder] concepts are based upon perceptual data, and higher-level concepts are based upon other concepts, which then in turn are based on perceptual data, are there... is there a range of optionality here in terms of concepts that might be formed either way?

GS: Uh, that... the question is, "Can there be concepts that are formed either directly from perception or based on concepts, that is, is there some concept which one might form just from perception, but, uh, one might also form it from other concepts?" Um... no. I don't think there are. Um... it's... it- it's first-level if it *could* be formed directly from perception--that is, if it's possible to see the relationship between the units on the basis of perception directly such that you don't need intervening concepts to do that. If that's the case then it's a first-level concept, uh... though in your own history, you might have... the way in which you formed it might have involved an earlier concept. So, for example, take, um, the concept, oh... in a lecture... oh what's... what's an obscure an- take the concept "aardvark." I'm not sure that I've ever seen an aardvark. Um... what I know about "aardvark" is it's a certain type of animal. So my particular "aardvark" knowledge might, in a way, depend on my animal knowledge 'cause I don't know anything about aardvarks, I've just sort of learned about it at... you know, from books where they listed animals, or say that was the case. Now, in fact the, the concept's a little bit floating for me, so I don't even think I really fully have it, but maybe there's some case where you could, um... you'd count as having it and you'd be able to recognize one if you saw one in some way, but in fact the way you learned about it was sort of round-about. Perhaps that's the case with, um... well, let's just stop at that o-... uh, I was gonna introduce a complication that would be confusing. Um... in that case "aardvark" remains a first-level concept; I'm just weird. That is, somebody who had a... somebody *can* form the concept "aardvark" directly from perception. If I

started seeing some aardvarks--I'd get into a place where I can do that--the concept "aardvark" does not depend on this strange circumlocution I may have taken in starting to have aardvark thoughts.

Um... the, the complication had to do with, "what if you discover an animal paleontologically," you know, by its fossils, which, if it existed now, you'd be able to form a first-level concept of, but I think the fact that you have to discover those concep-... those animals in that way means that they fall into a different category... uh, for us.

[someone is talking in the background but much too faintly to make out at all]

GS: the fossil itself is a perception, certainly. The question is, um, "Is, uh, stegosaurus a first-level concept?" You know, be-... we only know of the existence of stegosauruses through this complex, inferential process, though were one of them, uh, to be in the room, we'd sure... perceive it [laughter] and, uh, and be able to notice that it was a lot different from the rest of us.

Anyway, we're, uh, really out of time and, I think, a little over now, so we have to stop. [applause beginning on "so"]

Lecture 4: Objectivity

Yesterday we discussed the process of conceptualizing as a whole: the process of forming new concepts, either directly from perception or with the aid of former concepts, the process of applying those concepts, either directly on the basis of perception or through middle terms in inference, be that inference deductive or inductive. We spoke about the process, in the previous day, of defining concepts a little bit, and of changing those definitions as one's context of knowledge expands. This is the process of conceptualizing, and it is a complex, ongoing process. Throughout all this, the point that was in the background that this is a process over which we have conscious control, and this point now is going to shift into the foreground. Conceptualization is a process that needs to be enacted and sustained consciously. In order to know the world at the conceptual level, in order to conceptualize it, we need to act in a specific, complex way—in the way we've been discussing. We need to do this by choice. And thus we need a method for how to conduct our mental lives. We need to make the process of conceptualizing self-conscious. Now, the content of this method is basically what we've been discussing all along. It's simply doing the acts that we discussed over the past two days. But we've been discussing these actions not as a method but as a process—as a complex set of actions, but not as a complex set of actions that must be volitionally performed in a certain way.

4.1. Logic

The method, viewed as a method, is logic: "the art of non-contradictory reasoning," as Ayn Rand defined it. Logic, as a science, as a branch of study, is the study of the method as a skill. It formulates cannons of reasoning, rules of inference, rules of definition. We spoke about what it is to know something deductively by applying concepts, what it is to know something inductively, and of definitions and what their purpose is, though only briefly. Logic gives you specific standards by which you can evaluate inferences or definitions or concepts that are before you, and they're formulated so as to aid you in guiding your thinking as you sort of go through it. Epistemology studies methodology more abstractly. Logic tells you, for example, how to identify something's genus, or tell whether a given deduction is valid. Epistemology identifies the purpose and nature of definition and inference as such and their role in knowing as a whole. In this way, epistemology is more abstract. There's another sense, which you can ask about if you're interested, in which logic is more abstract, but that needn't concern us here.

This is not a class on logic, so we're not going to go through sort of in detail how one identifies the right genus for something, or so forth. But I do want to talk about the process that we've discussed yesterday and over the previous days from the perspective of it being something one has to choose to do, from the perspective of norms. That the process, as we've described it, of conceptualizing is not something happens but something that one does, and one can be assessed, and one can assess oneself, and needs to, as to whether or not one is performing that process, as opposed to some superficially similar deviation from it. This is the context for the introduction of the concept "objectivity".

4.2. Objectivity

4.2.1. Objectivity as volitional adherence to reality by the method of logic

"Objectivity" is the name for the norm of being logical. This is the sense of it that's used... "being logical" is too narrow, actually, or doesn't quite capture the right sense of it. It's the norm of being methodical where the method is true and based on human nature. This is the sense of the word "objective" as used in such phrases as "objective reporting" or "objective grading." I want to mention

another sense in which Ayn Rand uses it, and which is related, just to avoid confusion. She speaks often of "objective reality". Objective reality means (this is my definition): reality considered as the mind-independent object of consciousness which therefore sets the standards for cognition. Objectivity, in the sense we're discussing it, is methodical adherence to those standards. Thus Dr. Peikoff defines objectivity as "volitional adherence to reality by the method of logic".

4.2.2. The objectivity of mental products

"Objective" also denotes a category of mental products contrasted from subjective mental products, in that the objective ones were formed by the correct method. Objective mental products—objective concepts, objective conclusions—are the results of volitional, logical processing of perceptual data. They're the results of reality-based, logical processing of perceptual data. By contrast, subjective mental products are the results of non-reality based processing of perceptual data. All data's perceptual ultimately. So, we have a contrast between objective and subjective mental products—between objective grading and subjective grading, objective reporting and subjective reporting, objective reaching of conclusions in general and subjective reaching of conclusions in general—objective when it's reached by a proper method, subjective when it's reached regardless of the proper method.

4.2.3. Intrinsicism and Subjectivism

4.2.3.1. Intrinsicism and subjectivism as theories of knowledge

Two more concepts we need to introduce: "intrinsicism" and "subjectivism". We've spoken about objective and subjective mental products. Intrinsicism and subjectivism, and then objectivism, are theories of knowledge. Intrinsicism as a theory of knowledge is the theory that what are in fact mental products—concepts and conclusions, namely, and in general conceptual knowledge—exist independent of the mind, and, especially importantly, independent of our own processing—that they're some sort of objects in reality as such. The subjective theory of mental products acknowledges that mental products are the result of our processing, but thinks that this processing is necessarily non-reality-based—that there's no such thing as reality-based processing—and thus, that these products are not awarenesses, or cannot constitute awarenesses, of a mind-independent world. Intrinsicism is the theory that we are conscious no-how (or, at least, no volitional how); Subjectivism is the theory that we are not conscious, that is, that we're not aware of a world outside ourselves, but that we're not conscious somehow. There's a "how"—there's a thing that our consciousness is doing—but it's not "being conscious."

4.2.3.2. Intrinsicism and subjectivism as methods

Now, I want to talk about these not as theories of knowledge—in which one might just say, "knowledge is intrinsic" or "knowledge is subjective"—but as methods. There is no such thing as an intrinsic mental product in the world. There is such a thing as a subjective mental product, it's just bad. Intrinsicism as a theory of knowledge is the view that there are intrinsic mental products, and that things like concepts and conclusions are amongst them. Subjectivism is the view that all mental products are subjective, necessarily. I want to discuss these, though, as methods. Though there are no intrinsic mental products, there are intrinsic*ist* mental products. That is, there are mental products that are formed on a method, which method is dictated by the false theory of intrinsicism. And likewise, in addition to mental products that are merely subjective—that is, not based on reality—there are subjectivist mental products, that is, mental products that are based on a method, which method dictates

not basing them on reality, dictates just coming up with them sort of without reference to an outside world. As a cognitive method, intrinsicism consists in regarding one's mental products as unchallengeable... absolutes not in need of validation or of possible qualification in light of new knowledge. It's treating them as though they were revelations. Subjectivism is treating of all mental products as an optional matter. These methods are quite common.

4.3. Integration and reduction as the essentials of an objective method

The proper method, of course, is objectivism. And what does that mean? It means treating one's mental products as the products of a process which one had to perform, and had to perform in a certain way. Now what is required to treat one's mental products in this way? To treat them as the results of the processing of perceptual data in a certain way, which has to be done in a certain way to constitute knowledge.

Well, I think there are two processes that are involved in treating one's knowledge as objective, and thereby in making them objective. I get this point from Dr. Peikoff; you can find it in Ayn Rand's writings once he pointed it out, but OPAR's treatment of this issue is really a revelation (in a different sense than I was using that word a moment ago). There are two processes that are involved in treating one's knowledge as objective, rather than as intrinsic or subjective, and proceeding objectively in acquiring new knowledge and in holding one's knowledge, rather than treating one's current mental products either as unchallengeable revelations, which you don't need to do any further work on, or else as things that come from your mind that don't need to bear any relationship to reality. These two are reduction and integration. These are the two tasks involved in being objective. Reduction consists in tracing back the steps by which one arrived at a particular product from perceptual data. Integration is the task of relating one's mental products to the whole of one's knowledge—relating a given mental product that one's considering to the whole of one's knowledge. And it's these two things that we need to do to validate any conclusion or any concept—to validate any cognition.

4.3.1. How the needs for reduction and integration each flow from the nature of conceptualization.

I want to talk a little bit about how each of these flows from the nature of conceptualization. Each of these requirements flows from the nature of conceptualization as we discussed it yesterday. Remember that knowledge is acquired by the stepwise integration of perceptual data. Each of these steps needs to be volitionally performed *and sustained*. The sustaining part is important. Remember, knowledge is a grasp, and a grasp is something active; it's something that one does so long as one has a grasp of the thing one's grasping. It's something you need to keep doing after initially forming it. We've seen this in the case of concepts especially: they're institutionalized perspectives or policies. To form a concept is to put oneself on the policy of regarding things a certain way, and that's something one has to keep doing, it's not something that's over and done with when one forms the concept. Moreover, we've seen that the nature of concepts requires us to revise the details of that policy—the policy of grouping—as we learn more; it requires us to change our definitions in some cases. Since the process is complex and ongoing, we need to check and control it, and to check and control the process, to manage it, we need to consciously know each stage. Thus the need for reduction, the tracing back of a concept or conclusion to perceptual data through all the intermediate steps. That's why we need reduction. Remember, your previous knowledge on which you're building new knowledge is not like a brick that you formed by integrating some clay, and now you can, you know, integrate it with other bricks by putting mortar, you know, between them and building a house. The earlier formed knowledge is like a brick in that it's something you can build on, but it's not like a brick in that it's something

that's done with once one has made it. One has to hold each brick together, so to speak, and here the metaphor breaks down. This is why it's a much better metaphor to think of it like a cell or some component of an organism the very existence of which requires continued action. If that's right, then, in thinking, we are *enacting* all of our concepts—all of the concepts that are involved in the complex judgments that we make. We're not using previously formed concepts in a way that doesn't require current action. To identify something, say as an electron (or as an evil person, or as a hero) requires a complex process of action, and one needs to know all the steps in that process: one needs to know how one is interrelating the different things that one interrelates using the concept. So you need to know, in the present, all the steps that are involved in the process. That's why we need *reduction*. Even if it was performed properly in the past—even if the steps were right, you need to *know presently* that they are right—you need to see the connection—see the relation back to perception, while you're using the knowledge.

Why, then, do we need to integrate our knowledge? Well, I think the first and simplest answer to why we need to integrate a cognition in order to validate it, is that integration provides a check on our current knowledge. By interrelating ideas, we discover any contractions that there might be, which would be signs of error. But just to say this would be to relegate integration's role to that of a mere corrective. Remember that integration is our means of expanding our knowledge in the first place. Integration is how we enable ourselves to know more and more and more and more. And the integration of a concept is inherently a forward-looking process. To form a concept is to commit to further integration, and the same is true (you can ask about how this is so if you like later, although some of it'll be covered) in the formation of a generalization. The formation of any conscious content that involves an integration is forward-looking: it involves a commitment to keep on interrelating. If one stops working to interrelate—if one stops integrating, one abandons the policy in which all of one's concepts consist—one undoes one's concepts. Remember, the units of a given concept are kept differentiated from the rest of existence by a dynamic process. A process that includes, for example, changing definitions when necessary. If one drops the policy of integration, one starts collecting discrete bits of disconnected knowledge. To do this is to disintegrate one's current concepts, to stop them from being concepts.

Integration, we can say, is forward-looking, reduction backwards-looking. In knowing, one needs to know where one came from and where one's going. It's a motion—it's an action one performs—and one can never stand still. If one does, one's not doing the process anymore—one is no longer engaged in *knowing*, and one's mental content is no longer knowledge.

4.3.2. The relation between integration and reduction

Now, in philosophy, it's typical to prize either reduction or integration, though usually not in quite those words, at the expense of the other. When reduction is stressed, it takes the form of out-of-context, disintegrated proofs showing how one thing follows from another in a way that's isolated from the rest of all knowledge. Certain of the empiricist philosophers were trying to do this. They were trying to reduce to something that they take as first level—maybe it's sensations—and show how things go back to that, but each chain sort of in a void from all the others. And these views have many problems with them; among them is that they don't have the right basis, that is, they're trying to reduce to sensations, but there are other problems in these views, and people can ask about them.

I think the more common thing to stress, at the expense of the other, is integration, especially now. It's very common to take integration as the essence of a proper methodology, and to ignore reduction altogether. The idea is that a proper methodology consists in working to make one's knowledge coherent, regardless of what it is that one's making coherent. I've already indicated how this is and how it follows from false theories of concepts on the second day when I discussed theories

of concepts. I want to take it up from a different angle now. The idea—this coherentist method—it goes by a bunch of different names in different circles. Sometimes it's called "dialectic", though that's meant a dozen different things in the history of philosophy. Sometimes it's called "reflective equilibrium". There are different names for this. It's the process of shuffling around, the ideas in your head, noticing that some of then don't fit together—raising contradictions or tensions or problems, or just registering an uneasyness about them—and then tinkering with them until these problems go away. Now, you might think, "something's right about this—you do need to take all your different ideas and see if they're consistent, see if you run up against a contradiction," and that is true, that's what integration says. But integration is precisely worthless without reduction. Why? Well, what do you do when you come up against a contradiction? In a sentence Ayn Rand made famous in Atlas Shrugged. Everyone? "Check your premises." Which means what? Reduce. Ask, "How did I get to these premises which contradict each other? Something must be wrong on the way to here if here is a contradiction." And so you go back down, back towards perception, see where the mistake was, and then correct the mistake and integrate further. That's not the approach that's taken by anyone I can think of who has a systematic view of methodology, and it's not the approach that's taken, in their actual work, by most philosophers. What they do is try to go forward somehow past the contradiction—to transcend it, synthesize it, get around it somehow—rather than going back to discover and correct its cause. What this amounts to in practice is coming up with some new slogan or bunch of words which are confused and vague enough that when you talk in terms of them, the contradiction is obscured for a while. For example, people who know current philosophy of mind know that it was supposed to, for a decade or two, be a big solution to talk about all the things "supervening" on one another, precisely because nobody knows what it means to supervene on something. And then about ten years later people realized they didn't know what it means to supervene, and so there was another contradiction, and some other senseless concept will be introduced to solve the problem. This is the same way Plato's forms got introduced and every other bad thing in the history of philosophy. So I just want to point that out. Often people see people stressing integration and know that it's an important thing, integration, and they get excited and thing they've found allies; most often the people are not allies, they're Hegels.

4.4. Validating concepts and definitions

Okay, we've been talking very abstractly about objectivity, about the need to reduce and the need to integrate, about the possible errors of Intrinsicism and Subjectivism, the false methods. I want to talk about this in some more concrete cases or types of cases. First, the case of validating concepts and definitions. Definitions, remember, are statements of the policy in which one's concepts consist. Or rather, a definition is a statement of the policy, an articulation of the policy, in which a concept consists.

4.4.1. Why definitions are necissary

A definition is necessary because you need to know the policy on which you're differentiating in order to keep the concept alive and tied to reality. In the cases of low-level concepts, this information can be held perceptually, but in higher-level cases you need a conceptual definition. Ayn Rand aptly identifies the dividing line—it occurs as soon as you utter words with the feel, "I kinda know what I mean." To validate a concept and to validate its definition amount to the same thing. The concept is differentiating and integrating things on a principle; the definition is a statement of the principle. You validate the concept by showing that the definition is valid. In general, what we do in validating a concept or definition is show that the definition distinguishes units that are essentially similar to one another as opposed to everything else, that are similar to one another in a way that gives rise to many

other similarities to one another, as opposed to to other things. In the case of first-level concepts, and maybe in some other cases, but especially in the case of first-level concepts where the validity can be grasped perceptually, the main issue in validating a definition is not validating the concept really, but rather figuring out which of the many traits of the unit is fundamental. But this isn't really the question that one has to ask in the case of most concepts and definitions one's trying to validate. Very rarely is the situation that everybody agrees on what the units are, so that it's clear and uncontroversial and easy to tell what the units are and all one needs to do is state which of their characteristics is the essential, the fundamental, one. Usually it's unclear what things would fall under the concept at all—that's why one needs a definition. So the question that one needs to ask, really, to validate a definition is not, "Is this phrase, is this sentence the definition of this concept?" Rather, it's "Is there a need for a concept at all corresponding to this definition?" If the answer's yes, you can then worry about whether the word that you've got for it is the right word for it, and there are principles for determining that, I'm not going to discuss those, but people can ask in the question period. So the real issue is: Do we need a concept corresponding to this definition at all?

4.4.2. Methofological intrinsicism about concepts and definitions

And it's in this context that I want to talk a little bit about intrinsicism about concepts before saying more about the right way to do it, because I think this is a pervasive error made honestly and easily by people who are trying to be objective but who don't fully understand what it means in this sphere. I find that a lot of arguments—for example, on HBL—reduce to one or more parties making this error. The error consists in treating one's way of grouping in one's system of classification as sacrosanct and as prior to the definition, and then defining by simply putting the words to the policy that one's using, without acknowledging the need to justify or validate one's way of classifying. It amounts to taking your concept for granted and then documenting it rather than showing why you need the concept at all, which is what you need to do in validating a definition.

Now let me give you two examples of arguments that you're probably familiar with that make this error—or at least, as they're often put forward they make this error. Here's the first: "Of course Objectivism is a form of libertarianism. Libertarianism is the view that governments should do nothing but police people and run the military, and Objectivism holds that the government should do nothing but police people and run the military, therefore Objectivism is a form of Libertarianism. QED. Now yes, sure, I know you guys are different from the other libertarians, and I don't mean to deny that, but you can't deny that you're a form of libertarianism. After all, I just told you, libertarianism means xyz and you xyz, so you're a sort of libertarian." Okay, here's another one: "There's no such thing as gay marriage. There just can't be. Marriage is a relationship between one man and one woman and 'gay marriage'—it's like *two* men. Marriage means one man, one woman." Now the point isn't that the conclusions of these arguments are false. The conclusion of the first argument is false, I think the conclusion of the second argument is false as well, although I think that would be more controversial in this context. But the point is not that the conclusions are false, it's that there's something wrong with the way of arguing, even if the conclusion's right in one of the cases. It's okay to deduce from definitions in some cases, we've seen that that's sort of implicit in recognizing something as an instance of a concept, but that's only the last step, and it's never the controversial or difficult step. It's rarely the case that somebody acknowledges, for example, that marriage is a relationship between a man and a woman and simply fails to notice that gay marriage doesn't fit that definition, so that the person can be convinced by pointing that out. The problem is that it's not clear whether marriage *should* be defined that way or whether libertariasnism should be defined that way or whether it should be a concept at all. The problem with the intrinsicist is that he takes his definition as an unimpeachable beginning to knowledge. He doesn't really have the idea of validating his concepts, of identifying the

grounds for the policy on which he's grouping. And this is really what one has to do when validating a concept or a definition.

4.4.3. How to validate a definition

Now, how do you do it? Well, you have to think about the referents and think about why you group them the way you do and what things you need to know about them in order to see why they are to be grouped that way. That is, the grouping have to be based on an essential similarity between the units, on a similarity between the units as opposed to other things which explains many other similarities between the units as opposed to their foils. To validate a definition is to show how the units are essentially similar, and that often requires often, many intermediate steps because most concepts are not first-level, and most of the concepts over which people argue are very high-level. You can't just see how Objectivists are similar or different from libertarians. They're all men, you can see that at the first-level, and you need a lot more knowledge to have any idea about how to relate them to each other, and all that further knowledge has to be in play when one is validating one's definition.

Now, to say more about how to do it would be really to go into a clinic in logic or a clinic in defining, and that's not what I'm going to do here, but I want to point you to the best example that I know of how to proceed in this, and that's Ayn Rand's discussion of how to define "justice" in the chapter on definitions in ITOE. And notice that what she does is she doesn't say, "This is just and that's unjust. Now what do they have in common?" She says, "Why do I need a concept of justice?" And she goes through a whole bunch of stages of different things that need to be differentiated from one another. to form the concept. Here's a question I find helpful in getting myself into the right context—into the right mental set—to validate a definition. It might be a matter of my own psycho-epistemology, but it mighthelp you, too. I ask: "Why am I in the business of treating these things in this way?"

Now, I just want to keep in mind the role of reduction and integration here. This "How did I get into the business? Why am I in the business of grouping things in this way?" is the reduction of a concept. It's asking how the concept "came up." What things did I see, what path did I travel from perception that leads me to be grouping this way rather than some other way? But you also need to integrate. Why? You need to keep in mind the whole context. Whether a concept is needed depends on everything one knows, because the concept is meant to differentiate certain existents from all other existents, and, as one discovers more things, one may need to refine one's definition, or one may find that the concept wasn't proper, though it seemed to be in the first place—say a concept like "bilious," which grouped types of people based on whether they have a lot of bile (as opposed to pglegmatic people who had more phlem and so forth—these are early medical conceopts.) So one needs to always have sort of all of one's knowledge available to one when one's forming a concept and when one's validating a concept, and one can do that through keeping in mind the axioms and keeping in mind the wide concepts that are relevant to the area of life in which one is thinking, so that if one's thinking about any concept involving human beings, for example, the division between volitional action and automatic action.

4.4.4. Subjectivism about concepts and definitions

We've spoken about intrinsicism about concepts and their validation, as opposed to the proper view. I want to mention briefly subjectivism about concepts, which I think is also endemic—less so amongst Objectivists because once you grasp the need for objectivity it's harder to accidentally fall into it than it is to fall into intrinsicism. The idea of subjectivism about concepts is that way one classifies is arbitrary. One could just as well have classified some other way; it's a matter of preference. Mottos of

subjectivists about definitions: “Set up your definitions however you like, and then we’ll talk.” People will say that in arguments, you know, “Just define the terms however you want, and we’ll talk.” (It’s like the guy in *Atlas* who thinks he can live under any political system; he dies in the tunnel catastrophe.) Motto two: “I can’t *argue* for my definition, but try thinking of things this way, and I think you’ll like it too.”

Now this last motto is not really that funny. I mean, it’s not ridiculous, and people, in fact, do take more or less this—this view all the time. “If you think about things in these categories, it’ll make the world feel more orderly and sensible to you. It’ll help you; you’ll feel a sense that you’re making sense of things you couldn’t.” One of the dominant views of how science progresses is that people go through periods during a scientific revolution, in which the proponents of the new view aren’t able to *prove* to anybody that their way is better or right, but they say, “Look, start thinking about things my way, start thinking about things this way, on this paradigm,” to use the buzzword for this, “and you’ll see that it fits together well and makes sense. You won’t see *that* the new theory is right, but the world will seem right to you, it’ll seem to be a good and sensible place.” There’s a professor at Pitt, one of the better professors at Pitt—that’s the University of Pittsburgh where I study—who’s been heard to say, “Philosophers don’t argue, they paint pictures.” What does he mean? He means that they don’t try to prove, step by step, that one view is right and another is wrong. Rather, they sketch a whole way of looking at the world and compare it to other whole ways of looking at the world. Now, this is wrong, and it’s a bad view, but again, it’s not dumb—not if you think the alternative of going step by step would be somehow trying to prove, from “here, now, white,” like the Empiricists think... well, let’s not put it that way. Not if you don’t have a sensible, coherent view of how a step by step proof could work, and the people who oppose these philosophers don’t have really coherent, step by step view of how a proof can work. But, anyway, that is subjectivism about concepts, and it is an error.

The right view is that you can show that a concept is valid by showing that the units are essentially similar, that the definition of the concept distinguishes essentially similar things—things that are similar in a way that causes many other similarities. And showing that will often require the use of many other concepts. It will require many, many steps from the perceptual level.

4.4.5. Invalid concepts and uses of concepts

I want to mention, briefly, and really just mention, several types of invalid concepts. Package-deals—that’s grouping by non-essential similarities. There’s something similar about the entities, like there’s something similar about all “greenies,” but it’s not the kind of similarity that gives rise to many other similarities such that a concept is warranted. Ayn Rand discusses any number of package deals—I haven’t counted them, but quite a lot—in her writings. Another type of invalid, mistaken concept is a grouping of nonexistent objects. That is, the problem isn’t that the objects are misgrouped, it’s that they don’t exist to be grouped: “gods,” “duties.” Or the groupings of things based on nonexistent characteristics: “phlegmatic,” you know, another one of these ancient medical groupings that I don’t think is based on anything, at least as far as I know. These, I think, are the only two essential types of invalid concepts.

The other things that one might think of as invalid concepts are gross misuses of concepts: stolen concepts (that is, concepts used in disregard of their place in the hierarchy of knowledge), anti-concepts (invalid concepts of any stripe which are being used to replace and to obliterate valid concepts—that is they’re being introduced into discussion intentionally to obfuscate some distinction that’s important), floating abstractions (concepts which are ill-defined so that the person who’s using them is unclear what they mean), and frozen abstractions. (A frozen abstraction is a concept which one uses but one substitutes one unit of the concept for the concept as a whole. The concept is valid, an example would be the concept of “morality,” but when one uses the concept, one thinks not of

moralities in general, including any morality, but of some particular morality, say altruism.) I don't really have anything I want to say, in this context, about them, but I thought I should mention each of these as different types of invalid concepts or uses of concepts that Ayn Rand identifies in the course of her writing.

4.5. Validating judgments

We've been talking about validating concepts and definitions. We're turning now to a new subject: validating judgments. And in particular, I want to talk about hypothetical reasoning. But a little bit about validating judgments in general first.

4.5.1. The basis for judgments

In a judgment, the subject is identified by subsuming it, or some characteristic, as a unit of the predicate. Thus, in "Dogs are animals," we identify dogs by subsuming them under the concept "animal," and in "Dogs bark," we partially identify dogs by subsuming their action under the concept "barking." At the first level, this act, as we've seen, can be performed directly from perception. You can just notice that a given dog falls within the dog range. At the higher levels, as we've seen, inference is involved. We've seen how "this bird is an organism" involves an implicit inference from "This is a bird, bird is an animal," and so forth. "This bird performs cellular respiration" involves even more premises. Of course, the inferences need to be reduced back to first-level judgments. And the judgments need to be integrated with all of one's other judgments to validate these conclusions.

4.5.2. Hypothetical reasoning

But I want to talk about a special case: the case of hypothetical reasoning. Because a lot of our reasoning—complex reasoning—consists in hypothetical reasoning, and it introduces quite a few new issues.

4.5.2.1. The process of hypothetical reasoning

I want to give a very simple example of hypothetical reasoning. I was at a Halloween party once and there was a box—it was kind of like a shoe box) with a hole cut in the side of it, and there was a bowl placed in the shoe box with some kind of gross-feeling stuff in it. I forget if it was cottage cheese or peeled grapes—actually there were several shoe boxes like this, so they each had something else in them. You couldn't see what was in it because the hole was too small to look into and it was dark, but you can put your hand in and touch the stuff, and then you were supposed to guess what it was, and you were probably supposed to guess something ghoulish like eyeballs, but in fact, you know, you could try to get the right answer and guess what it was. Well, you can't immediately tell by touch what the thing is; you can't tell cottage cheese from everything else by touch or peeled grapes from everything else by touch. So you don't have enough information to make a judgment as to what it is you're perceiving—identify the thing, when you just put your hand in the "mystery feel box". you have to think about it, "round... squishy... soft... could be an eyeball, but I know they wouldn't put an eyeball in a in a shoebox; what else is soft and round an squishy, uh maybe it's some kind of fruit... well, cantaloupe is firmer, and..." You see what you do. Ideas come up to you for things that it *might* be, based on what you do know about it—e.g. it's texture and shape—but you're not sure that these ideas are right; you're not ready to judge that it's, say, eyeballs or grapes, so you draw on your other knowledge about each of these things that suggests itself to you that it might be to try to figure out which one it is, and you go

through a process of asking questions. Maybe you determine some experiments you can perform: "I'll see if I can squish it," you know. And that's how you tell.

Let's take a more dignified and pronounced example: the example of a differential diagnosis in medicine, which people like me who watch *House* will be familiar with in dramatized form. What happens? A patient comes into the doctor or to the hospital with a number of symptoms. I have a handout which I'm not going to go through in detail, but that shows a simplified case of differential diagnosis. It's Exhibit D on your handout. The patient comes into the doctor with some symptoms, say he has chest pain and trouble breathing. What does the doctor do? The doctor generates a list of hypotheses, several things that he might identify as the cause. The question that the doctor has to answer is "What is the cause of the symptoms—of the chest pain and the trouble breathing?" just like what you have to identify in the "mystery feel box" case is, "What is the weird feeling thing?" The doctor generates a list of hypothesized causes: maybe it's an infectious condition, maybe it's ischemic, etc., then he uses his knowledge of the candidates to gather evidence, the evidence needed to narrow down the possibilities and figure out which one is the actual cause. This is hypothetical reasoning. He poses hypotheses, uses background knowledge, which suggests the hypotheses to him in the first place, to figure out which questions to ask to determine which hypothesis is true.

4.5.2.2. The role of he hierarchy of generality in hypothetical reasoning

Now, I want to use this example and work with it for quite a bit. There's one that I want to notice about it right at the outset, before we turn to the main topic for which I brought it up. Notice that the causes that the doctor considers in this case are conceptualized at a fairly general level. If we look on the list of possible causes, grouped by type, we have infectious, ischemic, trauma, gastro-intestinal causes—all possible broad types of cause of the relevant symptoms. And under each of these types we have a number of specific causes, one cannot investigate two million possible causes all at once, or even forty-five. In order to function, human beings need to integrate. In order to hold and deal with all this information, we need to integrate it. And so it makes sense that one would have to deal in generalized categories first, either by broad types, or perhaps by body parts that could lead to certain problems, so maybe it's something with the heart, maybe it's something with the lungs, etc., and one will see them always doing this in *House*—it's fun to watch. So I want to just point out that one initially poses hypotheses at a fairly general level, and, as one progresses, not only does one pick one of the hypotheses, but one also further specifies it—the signs are that it's infectious disease, well, which one?

4.5.2.3. Evidentiary statuses (possible, probable, and certain)

The main thing I want to illustrate with this example of differential diagnosis is three concepts which we can call "evidentiary status concepts": the concepts "possible," "probable," and "certain." You can think of what one does in forming a hypothesis and then in reasoning with it like this: The hypothesis is a judgment held in abeyance. One is not ready to judge, in the "mystery feel box" case, "This is an eyeball," or "This is a peeled grape." One knows one doesn't have the information necessary to make that judgment. But one needs to make some judgment, and one has some evidence on which to get going. So one proposes several judgments, holds them in abeyance, and judges between the judgments, and makes judgments about what one needs to do to determine which judgment (if any) is the one to endorse. We have concepts that denote the status of judgments held in abeyance during this process.

The first concept is "possible." Something is possible when there is some specific evidence for it, and we'll discuss what that means later in more detail, but an example would be that it's one of a

few known possible causes for the symptoms exhibited by the patient in the medical situation or, perhaps, it's one of the few known wet, squishy things in the "mystery feel box" case. To hold something as possible is to think of it as a going option—something that one's considering as a judgment one might make. Now, it's here important to distinguish possibility in the sense I'm using it here (and I'm following Dr. Peikoff in doing this) from what we might call metaphysical possibility. Epistemic possibility is distinct from metaphysical possibility. Metaphysical possibility means the capability of an entity to act in a certain way. Eyeballs, for example, can be removed from the head and put in a glass; it's possible for them to undergo this—they don't continue to be eyeballs for very long, but they can be taken out of the head. That's a sort of gruesome example, but it comes from the Halloween context. Men can jump, etc. It's something they can do, it's possible for a man to jump. That, again is metaphysical possibility. Epistemological possibility is the status of a judgment which one has reason to consider but which one is not in a position to actually make. It is one of several options that it's reasonable to consider, that's what it is to be possible, in some context—reasonable to consider in some context.

The next status is "probable." Probable denotes a sort of going favorite in the process of hypothetical reasoning. Something's probable when the preponderance of evidence is for it. When one would have to act on it if one had to act on one of the options now. The third status is "certain." Certain is the evidentiary status of something that one has enough evidence for to render conclusive judgment. To conceptualize something as certain is to say "case closed," to endorse the judgment of it. So these are three statuses that take place within this context of hypothetical reasoning.

4.5.3. The arbitrary

4.5.3.1. The arbitrary contrasted with the possible.

There's a fourth status that is different from these three [viz. possible, probable and certain], and it needs to be contrasted with them—they need to be taken as a unit compared to it—and that's the status "arbitrary." The arbitrary is that which is asserted with no specific evidence. The status with which arbitrary might be confused is possible, so we want to distinguish them. It is possible that the person in our medical example has an infectious disease because he has chest pain and trouble breathing, and there are several infectious diseases which would cause chest pain and trouble breathing. It is not possible that the man has, at this point in the process—when the man has just walked in with trouble breathing and so forth—that he has some specific bacteria. That is, if you were to say of this man who walks in with trouble breathing, "Maybe you have an infection, with baceria X", where X is responsible for some small percentage of infectious diseases, which are responsible for the relivant symptoms in some percentage of cases, so that it's, say, a one in a million chance that bacteria X is what the man has. Or even a one in five-thousand, or a probably a one in one-thousand. In this case it is not possible. There is no specific evidence for the diagnosis that X is the cause that's operating in this case. It's an arbitrary assertion.

A more clear case of an arbitrary assertion would be, "Maybe you have cancer." And I just point to somebody who I have no particular reason to think has cancer. It's possible for anyone to get cancer, it's a thing that human beings can get, and infection with a certain disease is a thing that human beings can get also, and it's also something that can cause cough or chest pain and trouble breathing. But it is not possible in this context that the person has cancer, or that the person who came into the hospital has this particular bacteria, because there's no evidence for that in particular.

Now I have an example that makes this a little bit clearer. Suppose that a man drops dead in Bangladesh, and you are the first person on the scene, and you see that he's got a knife in his back. You can conclude that the guy was stabbed. He's lying on the ground with a knife in his back, covered with

blood. And you can conclude that he was stabbed and killed by a human being, since only human beings kill people by putting knives in their backs. That does not, however, make it possible that Tom Bowden did it or that Betsy Speicher did it or that Don Watkins did it, none... neither of whom you know anything about other than that they're human beings who can (therefore) stab people. Now, granted there are only, what, six billion people, so it's a one in six billion shot, you might say, that each of them did it. But that's ridiculous. There's no specific evidence for the person. The metaphysical possibility that they're one of the things that could have done this does not give any reason to think that they might have done it. To treat their having done it as a possibility, you need specific evidence—that is, you need evidence that makes it possible for you to perform the process of hypothetical reasoning with this as one of the options. You need enough evidence for it that it can be one of a small number of options that you can contrast with each other and pursue in finding evidence for it. And in a case of something like medical diagnosis where there are thousands and thousands or God knows how many bacteria and viruses and so forth, you do this by grouping them into groups. It's not possible yet that this particular bacteria did it, but it is possible that the person has an infectious disease, which then would have to be either bacterial or viral, and at a certain point it's possible that the person has bacterial pneumonia before it yet is possible that it's this particular strain. Okay. So. The arbitrary is different from the possible. The arbitrary is that for which there is no specific evidence, and this makes it essentially different from that for which there is some specific evidence because that for which there is some specific evidence can play a certain role in a reasoning process—the reasoning process that gives rise to all these concepts—whereas that for which there is no specific evidence cannot play that role.

4.5.3.2. Why an arbitrary assertion is neither a judgment not a proposition

Two more points about the arbitrary. The arbitrary is not a judgment. The assertions "Maybe the guy has cancer" or "Don Watkins killed the guy in Bangladesh," are not judgments. Why? Because there's no process of identifying, there's no act of figuring out how this relates to the other things taking place at all. One just says something, one just comes with something out of the blue. What one's doing is essentially different from what one does when one judges. Moreover, the contents of such an arbitrary assertion are not a proposition and therefore are not either true or false. The contents of an arbitrary assertion don't have any meaning because concepts and the propositions which are composed out of them are only meaningful as parts of a complex process of reasoning—of judging, of relating things to one anther. When the concepts are taken out the context of this process, they stop being the concepts that they are, and therefore they stop having any meaning. This is a process that one performs as a whole, not in isolated episodes. One doesn't just form the concept "man" and then apply it, with one's other concepts having nothing to do with the process. When one starts committing conceptual malfeasance—when one starts making things up out of whole cloth—one is performing an essentially different process from the process of conceptualizing, the process of which the formation and use of concepts are a part, and insofar as one's operating in this essentially different way, one is not thinking at all. One's concepts are not concepts, one's utterances are not propositions; they're noises, associations; they're something else. Therefore they're cut off from the context that would enable you to say there's any relationship between them and reality—to say, in particular, that they're true and false, that is that they correspond to reality or that they contradict it.

Q&A

I think I'll pause very briefly here to take one or two questions on the immediately present material, that is, the last two topics: validating definitions or hypothetical reasoning, before going on,

and then we'll go on further and then have more general comments.

Uh, Paul? Then the back.

PH [fairly clear, but must turn up the volume a lot]: Yes, I wanna kind of ask about this specific evidence issue and... and the broad categories [as/ash??] for the narrow ones. For instance, this comes up all the time in my job: a woman feels a lump in her breast, and so the doctor will say, "well, we'll start with either benign or malignant" and then of course, it would be arbitrary to say, "Oh, it's going to be a __" and then you list a subtype of a breast cancer

GS: Mm-hmm

PH: So you're saying it's... we don't have specific evidence that it's a whatever-whatever cancer, but I don't think... I can't see how we have specific evidence that it's either benign or malignant yet [loud cough nearby] and, so [loud cough again] should-... shouldn't those be regarded as arbitrary, or- or you said there's possible, is there *some* specific evidence that it's one or the other, that's

GS[much louder than Paul]: Well,

PH: that's where I'm kind of getting...

GS[overlapping the rest of his words]: Now I'm not a doctor, so I don't know if this is... what I'm about to say is true, but something in pattern like this must be true. Uh, here's the evidence: there's a lump in the woman's breast, which means there's a kind of hard, dense thing where there's supposed to only be soft things, right? Okay. So there's some limited range of things which you know are capable of causing hard, dense material to be present in a breast. Um... let's say, if you went down to the lowest level of specificity, there are two million such things; it would not yet be the case that each of those two million would be a possibility. However, you know that there is a hard thing there where there normally wouldn't be, and that there has to be a cause of it. If you can group the 2 million things capable of causing it into some small number of units, some small number of abstract types of causes, which are such that you can compare them, and by so grouping them gather evidence to advance one hypothesis and, refute another one, and so forth, then each of the several broad classifications, which collectively, cover the things which are metaphysically able to create dense stuff in a breast, would count as an epistemic possibility, because I think the concept "epistemic possibility" has to be defined relative to that process of reasoning. Uh, that's how I understand it.

Um... I don't know your name

JP [soft]: My name's John, John Pachet [sp?]. Um... so just to summarize what you just said, would you say that metaphysical possibility doesn't imply epistemic possibility?

GS [much louder]: Yes, that's certainly true. Metaphysical possibility doesn't imply epistemic possibility. Now, I want to point out that in most of... and in many, many cases, um, hypothetical reasoning involves cause and effect reasoning. Most reasoning involves cause and effect reasoning, and certainly the differential diagnosis case involves cause and effect reasoning 'cause you're trying to find the cause of certain symptoms. But they are separable issues. Uh, you can have hypothetical reasoning even where it's not a cause and effect kind of case, and that was the point of the "feel box," you know, the Halloween trick where you reach in and you have to come up with a hypothesis as to what you're touching. It's not really a cause and effect issue there, um, except insofar as cause and effect is involved in everything. Um... uh, and metaphysical possibility doesn't really...th-...the- the confusion between metaphysical possibility and epistemic possibility: part of the reason why it comes up so much is because, um, both are involved in thinking about causality in different ways. That is, uh, me-... hypothetical reasoning is so often involved in cases of trying to figure out complex causes, and the way one does it is one thinks about what things *can* cause this, that is, what are the metaphysic-... what are the things that part of their nature is such that they can cause, say, a cough, or a lump appearing on an X-ray.

Um, there was someone el-... oh, uh, David.

David: Um... on- on the cadaver in Bangladesh, uh, you said, obviously you said, "It's not possible that... we can't say that it's possible that Don did it," but you also said it in the form of "it's not possible," and I think... aren't both arbit-... it's arbitrary, so you can't say either of those things.

GS [overlapping his next words]: Um, I wouldn't

David: You can say it's possible, so...

GS [cont.]: say it's impossible, uh, because, uh, not... What I mean in saying it's not possible is we're just... not we're-... if somebody says "it's possible" you'd say "no."

David [quietly in background]: it's not possible to say that it's possible

GS: Um, and not... well, no. Possible mean-... possible denotes a certain status--a certain evidentiary status. "That Don did it" does not have that evidentiary status, therefore it is not possible that Don did it. Uh... there I don't think is a corresponding sense of impossible, uh, in epistemic possibility. That's a metaphysical term. If you wanted to say, "it's impossible" in epistemic terms, it would mean "it contradicts something I already know." But that would really end up meaning it's metaphysically not possible for him to do it because, for example, say you know that Don was in Manhattan and you can't get from Bangladesh to Manhattan in time, and so forth, well then Don no longer metaphysically could've killed the guy since Don's not able to go faster than the fastest transportation we have, and he's not able to kill people in places where he's not present. At least, I hope--that's why I stay in Pittsburgh, I know he's...[laughter]

4.5.4. Concluding remarks on hypothetical reasoning

Just to sum up, I'm not really saying how, in this discussion of hypothetical reasoning, you prove the hypothesis that turns out to be certain. In pattern what you do is you show that the subject falls within the range that is the range of the predicate concept. That is, to show that "this disease is pneumonia" is to show that this disease has the distinguishing characteristic of pneumonia, it is like the other cases of pneumonia in precisely the respect that they're different from cases of non-pneumonia. How one does that is difficult, and to know how to do it in the case of pneumonia you'd need to be a physician. At the most general level, all epistemology says is that one does it by showing that it falls within the range of the concept. It's easy to see how that works in perceptual level cases. In higher cases that involve a lot of complex reasoning, you can see that it involves a lot of steps, which steps are reducible to the perceptual level cases, but to say what they are in any field, you'd need to be an expert in that field. What I'm doing in discussing hypothetical reasoning is identifying what it is to have a hypothesis and to reason in this way. It is to project judgments that one does not yet have the evidence to make, but which one has the evidence to know are one of a set of judgments which ought to be made and then to use that context to figure out which judgment in the set is the right one to make. You don't necessarily need to have an exhaustive knowledge of every possibility in the set, but it helps.

4.6. The status of validated cognition

4.6.1. Contextual absolutism

I want to close this positive presentation (and then we'll have questions again), by talking about the status of conclusions once one's validated them, once one's proved them, once one has reduced them to the perceptual level, seen how one reaches them from perception, and once one's sees how they relate to the rest of one's knowledge--that is, when one's integrated them. You need to integrate and reduce, reduce and integrate, in order to validate a conclusion. Such a conclusion has to be held as

absolutely true, as absolute but in the context that one is in, as a contextual absolute.

One does not know everything, and since everything is interrelated, one does not know everything about anything. That is, there is no thing about which one knows everything since to know everything about it would include knowing how it relates to every other thing, which one couldn't know without knowing everything as such. But this does not mean that one does not know the things that one does know. For example, to build an example Dr. Peikoff gives in *Objectivism: the Philosophy of Ayn Rand* about blood types: if you don't know the cause of every possible incompatibility between two quantities of blood, this does not mean that you do not know, for example, that the presence of a certain protein in any quantity of blood makes it incompatible with another protein in any other quantity of blood. That is, that A and B bloods are incompatible. And it does not prevent you from knowing that this A-B sort of incompatibility is the cause of all the known cases of blood incompatibility that have been studied. The case is that one has noticed some bloods being incompatible with others, such that when you perform a blood transfusion, one finds that a negative reaction occurs—I think a whole lot of clotting, although I'm not positive. One can discover that this is caused, in the cases one has seen, by the presence of two incompatible proteins in the different bloods, or rather, one blood has a protein for which the other blood has an antibody for, so it attacks it and forms clots. If you've discovered that and you've looked at a range of cases, you know that A blood—that is, the blood with the one protein—is incompatible with the B blood, and that this explanation of A-B blood incompatibility is a sufficient explanation of all the known cases of incompatibility. So those things one does know. And that qualifies as knowledge, even if one does not know what other factors could possibly come up in the future to lead to incompatibility between two quantities of blood. Not knowing everything does not invalidate the knowledge one has, unless, that is, one misformulates the knowledge one has so as to make it hostage to each new discovery. It's a difficult question regarding historical thinkers whether they made this mistake—that is, whether they said, in effect, "all A bloods will be compatible everywhere and anywhere, in any context, and it cannot possibly be the case that anything new will ever be discovered about bloods other than what we know now." If they held their view in that way, it would be a mistake. It's often difficult to know whether people made that mistake, in part because the context sets the meaning of the terms in which makes one's formulates one's conclusion. If a doctor says, "A bloods are compatible," for example, at a certain stage of knowledge, what does "compatible" mean? It might be defined, in terms of the absence of some interaction that occurs between A and B bloods. And that might be the proper way to define it in that context. Again, I don't know. To know this in any field would be to be an expert in that field. All one can do as an epistemologist is talk about what's true in pattern. But it's not important to determine about historical figures whether they held their conclusions in the right way. At least, that's not what's of primary importance. What's of importance is how we should hold our conclusions. And we should hold our conclusions by identifying what we know and what we do not know. Identifying, in the case of generalizations, that we've identified one causal factor, say, but that we don't know what other causal factors could be present. What one has proven, one's proven, and it's knowledge, but again, one does not know everything else that could be relevant to the situation at hand. These kinds of cases are difficult, and people can ask questions about them, we can talk about them more in the question period.

4.6.2. The possibility of error

I want to make one last point, and then we'll close. Neither the contextual nature of knowledge (that is, the fact that not knowing everything doesn't prevent you from knowing what you know) nor the method of objectivity makes you immune from error. If what you say is false, it's false. For example, even if a given doctor is right to think that a given quantity of blood is safe to give to a given patient, if he gives it to the patient and the patient dies as a result, the doctor's thought was wrong, even

if the doctor was perfectly objective, even if he was basing his knowledge on an induction which was validated and which was true (though it didn't take into account some factors that the doctor didn't know about and couldn't have known about, which happened to have interfered in this case). This is true even if the doctor, or whoever it was, was perfectly objective—even if he was certain that, for example, the blood was safe. I don't mean *felt* certain, but I mean properly judged himself to have enough evidence. You can be certain and wrong. Certainty is a judgment like any other judgment, and it takes place in a context. In the context in which you make the make the judgment, it can be the right judgment to make, and in that case you really were certain. Now that you've seen the patient die, for example in our medical case, you're no longer certain that that was the right thing to do. That is, you wouldn't do that for a similarly situated patient in the future. But you were certain at the time, and if your thinking at the time was unimpeachable—that is, it was perfectly logical and objective—, then it was unimpeachable—it was was perfectly logical and objective. Objectivity does not make you immune from error. What it does is make error an accident. It makes it an insignificant accident, something that is of no metaphysical importance, and which does not happen often. And it makes you progressively less susceptible to error with each new thing you learn. That's what objectivity does.

4.6.3 The role of objectivity in conceptualization

Now, I want to just differentiate this from a confusion that you can make. Objectivity is not just doing your best. It's not just using your mind to the fullest extent of your ability; that's rationality. Objectivity is following a method of cognition based on your knowledge of the nature of reality and on your knowledge of the nature of your consciousness. Objectivity is methodical rationality. It's being rational when you know how to be rational in conceptual terms. With thinking, it's performing the process of conceptualization on principle, in accordance with known principles of how to think.

Epistemology is the science that discovers those principles in the broadest terms. It discovers how one knows so that one can enact the process of knowing self-consciously. And we've discussed in outline how one knows over the course of the last two days, that is, the day before yesterday and yesterday and somewhat the day before. What we've discussed today is the perspective that one has to take on that process, given that one has to do it volitionally. And we've differentiated that perspective, we've differentiated the need to be objective from other ways that one might approach cognizing, one might approach thinking, trying to conduct one's mental life, if one had a false view of the nature of knowledge, or no view.

Q&A

Uh, Charles?

CK [soft but audible]: Yeah, could you distinguish between an epistemic certainty like you've described in the case with the blood--it turns out to be wrong, but, uh, the doctor was being objective, and- and it was...

GS [louder]: This is the case with a particular patient...

CK: Uh, yes

GS: ..., you know, "John will improve if I give him two cc's of this or whatever

Charles: Right, can you distinguish that from um, an axiom, which you wouldn't want to say, "well, it could turn out that you were wrong, that [he chuckles] "existence exists" wasn't

?female: yeah

GS: Can I distinguish a certainty from an a-... Well, first of all, certainty is a status that a proposition has. It's the status of saying "case closed" on it and making a judgment, so it's not the sort of thing that one would distinguish from an axiom first off. It's an epistemic status. Axiom is a- a thing

that is known immediately, that is, is known without any inference in perception, and which is contained in all other knowledge. It's presupposed by and contained in all items of knowledge. So they're not concepts that are directly on a CCD with one another. Um, now the point that you want to make I- I- the- the- I think the point that you're making with the question is that it's impossible for axioms to turn out wrong, uh, whereas it's possible for certainties to turn out wrong. Um..... in a way that's true, but in a way that's not true. I mean, we have to identify our axioms as axioms by following a certain method, and somebody can mistakenly think something's an axiom when it's not. Now, if it's really an axiom then it can't turn out to be wrong, but you can take it to be an axiom when it's not, and people do. So determining what things are your axioms, proving that it's an axiom, is, um, you know... that's part of objectivity and requires a method. And there are plenty of cases in the history of thought of things that were thought to be axioms and were not. Uh, the existence of god in some cases, for example. Um... and, uh, so forth.

Um... over there?

GB: My question has to do with uh... subjectivist mental products

GS: Uh-huh

GB: that's what you called it, um, and optional values

GS: okay

GB: And, um, I guess the thing is--I'm just thinking about values, just in general, and if you can consider them mental products, or if they're just... like I know the definition of a value is something that you act to gain or keep, so it's like a process, is it, but

GS: Um

GB: you'll also sometimes think of it as a

GS: Values... the question is, "Are values mental products, uh, are optional values subjective mental products or subjectivist worst... worse mental products?" Um...in the s-... well, one: values are somewhat beyond the scope of this course because this is uh, epistemology not ethics, but everything's interrelated. Uh, a value is a mental product in the sense that I'm using the term. Uh, in a certain sense there's no such thing as a mental product. That is, there's no, uh, thing that comes out of a mind and is... no act of consciousness is sort of finished in the way that, uh... and- and counts as a product of a mental process in the way that, say, a shoe counts as the product of a cobbler. He makes it and then it's, you know, just out there in the world independent of it. Um... all things in consciousness are active processes. Um, but what we mean by saying that something is a product is that the process has an enduring and unitary nature--it's something that exists over time. And we can use "product" more generally to refer to any mental state--like a conclusion, even that's not a conclusion that one holds over the course of years; it might be one that comes up in a particular situation and then one forgets about it. And that's the sense that I was using it in most of the discussion today. Now, values c-... are products in both senses. Values are ment-... in- in the... the less demanding sense, all values are products. The thing that's a product is not the object that's a value--that's a product in another sense: you have to produce it--but not a-... a conscious product in the sense we're using here. What's a product is, in effect, the valuing of it, right? Now... that is a mental product, or can be, in the more enduring sense of mental product, too. For example, your career ambition--that's a value that's something that-... that you've, uh, created. Um, you value that and you've institutionalized valuing that over time, and you're pursuing that. Your, um, your, you know, the love of your life, is- is similarly an institutionalized value in that way, analogous to a concept being an institutionalized perspective on things. And also moral... uh, you know, abstract values like, say, purpose is an abstract value, or self-esteem, um... so those are mental products. Now those can be subjective or objective. That is, you can form them in a reality-based way or in a non-reality-based way. And they can even be subjectivist or intrinsicist--that is, you can form your values on principle, and the-... You can form them in accordance

with a method, and the method can either be objective or it can be an intrinsicist method. You go about forming... You go about evaluating things in the way that you would have to were intrinsicism true, or in a subjectivist way--you regard all values as on a par and as a- a consequence of, you know, whatev-... of people's feelings. Um, optional values are something different altogether. An optional value is an objective value, um, which differs from person to person. That is, uh, it's objective, say, your valuing for your particular job, uh, it's objective that you need a productive purpose in your life, that's an objective value that everybody... and it's a universal objective value. Um, it's also objective for you to value your particular job, if it is--that is, if you've selected the job rationally. And there are standards for what it is to select a job rationally, and that would be another talk. Um... Or you can, um... but your job is not something that all men in general should want, or even that all men in your sort of circumstances should want. It's optional, it's something that you choose personally, but you can personally choose it in an objective way. Um... now this is a sort of complicated issue to get into all the details of it, because the kinds of things that are optional values, uh, and their relation to objective values, there are sort of a lot of subtleties there, and that is really for a course on ethics, um, so that's all I have to say on it now.

Uh, Harry?

HB: I wonder if you, in answer to Charles's question if you would... if you, um, agree or disagree with a point Dr. Peikoff has made about certainty not applying to things like direct perception because they're not... that certainty's reserved for the end stage of a process of accumulating evidence over time.

GS: Yeah, there's, um... Dr. Peikoff has made the point, and he makes it in- in OPAR, uh, also, that certainty doesn't apply, or at least doesn't primarily apply, to things that one knows directly. Rather, the concept comes up only in the context of, um, having to amass evidence for things over time, so largely in the context of hypothetical reasoning, and that's the context in which I discussed it. That's definitely true. Uh, we apply it to other things, um, because... they, like the things that are at the terminus of that, uh, process are things that we have a-... a "case closed" sort of judgment on, things that we are right to take as knowledge. Um, they have, in fact, a more secure status than the other things because there were less steps involved in getting to them and so less possibility of error--or they can have a more secure process--they are the things to which we will reduce the other things. Um, so they a... um, Since there's less possibility of error in the first place, and there's less that goes into forming them, uh, if we only had knowledge like that, we wouldn't need the concept "certainty." Certainly that's true. Um, I think it's appropriate to- to use th-... extend the concept "certain" to them. Um, and in the case of... of axioms, um, the things that are actually axioms are directly known, but, um... it is possible to confusedly thing something's an axiom when it's not, and when it's not even true, and so I think that's, uh, relevant and needs to- to be brought up in this context. I think that's all I- I have to say on that.

Uh, Betsy?

BS: Yeah, in your example where, uh, the doctor concludes that the patient will respond positively and doesn't, uh, if he's certain, is that really rationally warranted, or is he arbitrarily going beyond what he knows: that all treatments don't work for all people in all cir-...you know, of a certain type.

GS: Yeah, this is a good question. The question is ta- uh, I mean, that I find of interest. Take the doctor who concludes that the patient will respond well with this, uh, treatment, and the patient dies, um, and he dies as a result of that treatment because it turns out that there's something that he didn't know about the treatment or about the patient, and couldn't have known. Um... Betsy asks, "Is the

doctor really certain, or properly certain, in that situation." Well, let's... let's assume that the doctor behaved optimally rationally, and he did the right thing. That is, given that evidence and given that the guy was dying if he didn't do anything, he had to take some action, and this was the right action to take. I think we can agree on that, tha-... at least that there are cases like that. So the question arises, "What purs-... how should the doctor have viewed this? Should the doctor have viewed this as 'the best thing to do, although I'm not positive that it'll work'?" In some cases, certainly, people ought to view things like that; um, often we have to act when we know that we don't have enough information to be certain. Uh, the question is, "Can it ever be the case that the doctor reaches certainty, and is right to say "I'm certain this will work" when in fact it won't. Now, I don't kn-... have an opinion on medicine, that is, a lot depends on what the state of knowledge in medicine is, as to whether a doctor can ever be in that... uh, in that situation in medicine. I think, um... it seems preposterous to say that, um, you're not certain that it's a good thing to clean a cut and put a bandage over it. Um, I would think that would be certain, even knowing that there are lots of things about the body we don't know. But, again, I'm not an expert in- in medicine, but I- I think that's right. Uh, but going to cases other than ones in a scientific specialty where, by definition, if you're not in that specialty, you're not in a position to make the judgments, take the case, which, um, has been used as an example of this before, of, uh, Rearden's view of Francisco as, in effect, a hypocrite, which he forms after Fransisco doesn't stop the shipment of D'Anconia copper to Rearden being sunk when he knows how much Rearden is depending on it. Um... and after Rearden, uh... Fransisco has said so many inspiring things to him about, you know, productiveness and virtue and so forth. Um...

BS [quietly]: Was he certain there?

GS: I th-... well, that's the question. Um, I think, in that context, and in a such a context, one could be. That is, the kind of factors that are involved are so weird and unusual it would be cognitively stultifying to allow for their possibility. That is, it would be arbitrary to think, "Well, maybe this guy who seems wonderful is deceiving me as part of a deliberate act to try to destroy the whole system in which we're living, and it's led by this guy he met in college, and he's hiding out in a gulch half the--" [general laughter] I mean, it would be arbitrary by the setup of the thing. And so if one carried in one's mind even enough doubt to allow f-... for that possibility, one would have, in effect, reduced oneself to skepticism. Uh, one wouldn't be able to function. That is, one would be holding, always, an extra unit in one's mind that one would never have occasion to use. Once in your life, if you happen to be living in Atlas Shrugged, you know, the novel [Betsy seems to be saying something below you but can't be distinguished] you'd use it. Now where the cut-off is between that, uh, and where you only have a very high level of probability, and in what sciences certainty, rather than a high level of probability, is achieved, and at what point, is an issue for people who know a lot about those sciences to come up with. I think about cleaning a cut and putting a bandage on it. We almost certainly have it in medicine. Uh, do we ever have it about, uh, say, the results of a surgery or something like that? You'd have to ask a surgeon, uh, who was epistemically savvy.

Okay, um, Don?

Don: Um, I was hoping you could comment on the point you brought up yesterday about complex propositions.

GS: Oh, yeah. Um, that's a good question. Um, uh, when we've been talking about propositions, we've been talking about propositions of the form "S is P"--what are called predications or simple predications. You have a subject, say "man," and a predicate, say, um, "laughs" or "is rational" or "wears watches," and you just, uh, subsume the subject or an attribute or characteristic of the subject under the predicate. But there are, uh, more complex sorts of propositions or things that are thought to be propositions, and those get a lot of attention by logicians and people. Uh, that is, complexes of propositions made by logical connectives, things like "or" and "if." So, for example, um,

“If man is the rational animal then he requires freedom to survive.” That proposition. Um, now, I-... there are... have been a lot of attempts... I think the proper way to understand these types of propositions is to take very seriously the form-object distinction, uh, as is the case with understanding all propositions, and to think about... and- and the fact that consciousness, you know, is something that’s active, and to think about what it is one’s mind does when one has a thought of the form, um, “if man is rational then freedom is a requirement of his survival.” Or, um, “Either Fransisco is honest or dishonest”--that’s another kind of, you know, two propositions connected by a- an “or” in this case. Um... I think in this case the proper way to understand what one’s doing is not as asserting a conditional, that is, asserting, “if so-and-so then so-and-so,” which is how logicians usually think of it, but rather as asserting something conditionally. That is, the “if” part is a part of the form, not a part of the content of what one’s doing. Now, it has a content of its own, in a way, but it’s not part of what one’s asserting, so I don’t think it’s strictly true or false that, um, “if man is rational then he needs freedom.” Rather, it’s true that man needs freedom, and it’s true if he’s rational. So if you don’t know if man is rational, somehow, um, uh, and you were thinking about the question, you in effect give yourself the order--when you say, “If man is- is, uh, rational, does he need freedom?”--”Suppose man’s rational. Now, on that supposition, figure out whether he needs freedom.” And I think if one takes that attitude towards complex propositions it clears up a lot of issues in, uh, in logic, particularly contemporary logic, which is very confused.

Uh, one more point on this general theme, and we’ll close with that--notice how this relates to the issue of hypothetical reasoning. Uh, what one does when one says, “It’s probably cancer,” say, is one does not say... the content of what one’s saying is not “It’s probably cancer.” Rather, one says, “It’s cancer”, probabilistically. The “probably” modifies the saying, not what’s said. It’s a feature of the form of what one’s doing, it’s part of the identity of the act of consciousness, not part of the identity of the object. What one’s thinking about is its being cancer. What one’s asserting is its being cancer, but one’s not asserting its being cancer, one’s rather putting it forth proba-... as probably being cancer. Um, so I think that, um, relates back to the topic that we started the course with, or very early in: understanding and taking seriously what consciousness does--consciousness as an active process--and the acts of consciousness having an identity, having features that distinguish them from one another, and many of the problems in philosophy coming from a forgetting of that fact, of building the identity of the act into the identity of imaginary, special objects for consciousness. Um, that is, objects that are made up, you know, as a result of a false theory.

Okay, uh, that’s it for today, we’re out of time. Thank you. [applause]
[noise of mic being removed]