Concepts, Essences, and Mental Representation

Nathaniel D. Serio

A thesis submitted in partial fulfillment of the requirements for the degree of

Bachelor of Arts

from

Rutgers University

in

Philosophy; Cognitive Science; Linguistics

Kristen Syrett (co-supervisor)

Paul Pietroski (co-supervisor)

Contents

Abstract

Psychological essentialism is the view that certain concepts, more specifically, natural kind concepts like DOG, FEMALE, WATER, and so on, are constructed against the background of a intuitive folk belief that they have underlying "essences" which determine category identity and category membership. As is the case with most views, they have critics and psychological essentialism is not immune to this. Other studies have begun to call into question the strength that this proposed intuitive folk belief has on our categorization judgments (Malt, 1994; Tobia et al., 2020). This work also seems to be bolstered by more recent hypotheses about the relationship between lexical items, word meanings, and concepts, specifically, that it is not an isomorphic relationship. (Pietroski, 2017, 2018). The studies that I present in this work add to this line of criticism against psychological essentialism and, ultimately, these data argue that the previous results which had been chalked up as evidence in favor of the view are actually an effect of a certain experimental task setting. In this work I will present a way of accounting for these data on what I call the "one meaning–potentially many concepts" view. I also offer some potential directions for future research.

A primer for non-specialists

Let's think about thoughts for a moment. Among many things, I can think the following thought: "roses are red." Based on the fact that, one, we're all humans that have the ability to think and, two, that we're all probably familiar with the thought "roses are red," we are presumably in agreement that this thought, you know, whatever "thoughts" are, is about something in the world, namely, roses and how they are red. This "aboutness" property of thought is called *intentionality* – that is, thoughts, whatever they are, are things that can be about objects or events that are (i) not present to the thinker ("the apple in my backpack"), (ii) distant in time and space ("dinosaurs from the Mesozoic Era"), (iii) hypothetical – perhaps lying in the future ("the Chicago Bears might be in the playoffs this season"), or (iv) entirely imaginary ("Sherlock Holmes being a detective") (Shea, 2018).

So given what has just been said, here is a rough approximation of what a thought is: a mental representation with a particular content or, in other words, a particular "thing it's about." Ok, that's all fine and well, but what do we mean by this technical notion of *mental representation*? Well, based on the fact that we're all humans that are raised in linguistic communities and, as a result, we all typically develop some form of competence in a natural language, we are all familiar with how words and symbols can act as representations (e.g., $ represents currency). By analogy, we can think of thoughts as working in a similar fashion, it's just that with thoughts the representations occur "in the mind," so to speak. One way of elaborating this notion of mental representations then is to say the following: mental representations are psychological objects that are connected with meanings, you know, whatever they are. Naturally, however, it seems this way of putting things raises two follow-up questions…

(1) How do mental representations get their meanings?

(2) Since mental representations are associated with some particular content, what is the structure of that content?

The first question is, I think, fairly straightforward in terms of what it is asking, but the second question can be difficult to make sense of. To understand exactly what (2) is getting at, let's revisit our example involving the thought "roses are red." So far, we have stipulated that thoughts are mental representations associated with a particular content and so, for any given thought we might reasonably ask the question "what does the content of that mental representation look like?" For instance, with respect to "roses are red," it could be imagistic – i.e., we can imagine a red rose and see it with our mind's eye. Or it could be propositional – i.e., something possessing a "sentence-like structure" that highlights logical relations and truth conditions. Or maybe there is something to be said for both of these things. In any case, it's worth noting that (1) and (2) are two of the $64,000 questions in philosophy of mind and cognitive science.

The current project approaches the above questions from a slightly different angle. In particular, I am interested in *concepts*, which for our present purposes can be thought of as mental representations that are associated with lexical items or words. In keeping with the "roses are red" example, then, there is a particular mental representation that is associated with the lexical item 'rose' such that it allows us to think about roses and it is that thing which we are talking about when we talk about the concept ROSE.[1] Calling attention to the classic architectural metaphor, a thought can be understood as an edifice and concepts as the building blocks. Seeing it as that now we are beginning to put together a rough version of what gets called *The Representational Theory of Mind*,

[1] Small caps is technical notation for referring to concept associated with 'rose', where single quotes notes the lexical item.

I shall want to add that it isn't enough for a theory of mind to only tell a story about representation, rather, it also needs to tell a story about higher cognitive competences (e.g., categorization, induction, speech production, understanding, etc.) (Fodor, 1998).

As it happens, the opening line to such of story, at least in its contemporary formulation, claims that the mind is like a computer and as a result of this metaphor, higher cognitive competences are like "computations." On this way of viewing things, then, concepts are the *symbolic structures* which the computations perform over. In case this is confusing, let's consider a toy example. Suppose there is a machine which performs a simple computational procedure using the following rule:

(R) If A, then type B

Immediately, we should take notice that with respect to this rule, A and B are symbolic structures and moreover, that the machine is programmed to carry out a particular manipulation of these symbols under a certain condition, namely, "If A." So, if the following information served as input to the machine and it was successfully able to perform the computation, then this is what would occur as a result:

Input: A

Output: B

If we shift our attention back to concepts now, we shall want to recall that we described concepts as mental representations (just specifically those which are associated with lexical items or words) and that we described mental representations as being associated with some sort of content. All of which is in an effort to point out that concepts have content (called *conceptual content*), so, given that our proposed theory of mind analogizes concepts with symbols, we will also to take careful notice that

they are more they just "plain" symbols like A and B from the toy example. To elaborate, A and B are just arbitrary symbols that I chose at random – they don't have any content – concepts, however, do have content and thus, computational processes in the mind need to be sensitive to their constituent structure. That said, we might wish to follow Fodor (1998) and clarify this notion of a computation by saying "[it] is some kind of content-respecting causal relation among symbols" (p. 11).

To be sure, there is a lot more discussion that could be had here, but for the purposes of a primer section this should be enough to get the ball rolling. For now, we can consider the above to be an acceptable *pro tem* theory of concepts. But before we wrap up this section, there is one addendum I need (and really want) to make. The addendum is this: since we understand higher cognitive competences on the analogy of computational procedures that are supposedly content-respecting causal relations among concepts, we can make an important methodological inference for how to study concepts. More specifically, we can infer that if we study the outputs of these higher cognitive competences in certain experimental task settings, then those results will decidedly tell us something about the content of the concepts serving as inputs to those processes. This is an important point to keep in mind, because, one, it is a fairly ubiquitous methodological inference that gets made in cognitive science and, two, it is something which I make reference to repeatedly throughout this work. In fact, I would go so far as to say that all of experiments which I present below have been conducted against the background of that very methodological inference. Anyway, so much for a primer, we are now able to move onward to the main work.

Introduction

Psychological essentialism is the view that certain concepts, more specifically, natural kind concepts like DOG, FEMALE, WATER, and so on, are constructed against the background of a intuitive folk belief that they have underlying "essences" which determine category identity and category membership. *Prima facie*, this view might seem to bear a striking resemblance to Putnam (1975)'s view of word meaning, however, to those people who might feel inclined to make such a connection, I should caution that there are some noteworthy differences. For one thing, Putnam is making a quasi-metaphysical proposal about what meanings are and how they determine the extensions of words. Putnam thinks meanings are constituted by essence-like properties (e.g., Water = H_2O) and that, whatever meanings are, they "just ain't in the head!"

Of course, what the philosophers care about and what the psychologists care about are usually different things and so, the proponents of psychological essentialism (mostly psychologists) tend to be agnostic when it comes to metaphysical claims about the structure of reality or debates over how meanings are determined. On that note, I want to emphasize the following point: psychological essentialism is not the view the things in the world actually have essences or that meanings are constituted by essence-like properties. Rather, the view claims that our concepts for certain things in the world might reflect such beliefs (Ahn et al., 2001; Gelman, 2013; Medin and Ortony, 1989). The point is, I think, summarized perfectly by Gelman (2003) when she argues on behalf of psychological essentialism that "[m]eanings may not be in the head, but the conceptual underpinnings to such a system may imply a kind of essentialism (in the head)" (p. 16).

At this point, we might wonder what reasons we have for thinking our conceptual

underpinnings "imply a kind of essentialism," to use Gelman's words. As it happens, support for psychological essentialism comes from (or rather, is said to come from) several sources. One major piece of evidence comes from Keil (1989), in which participants were asked to provide category identity judgments in response to vignettes of a natural kind or artifact undergoing some sort of transformation. In the vignettes, the transformation was always described as resulting in some of the thing's properties being altered, more specifically, as always altering the perceptual or superficial properties but not affecting the causal or essence-like property. Keil found that participants were more likely to resist judging that the thing had undergone a change in category identity when the thing under consideration was a natural kind as opposed to an artifact (e.g., people judged a raccoon would not transform into a skunk but a teapot would transform into a birdfeeder). In line with the predictions of psychological essentialism, then, these findings seem to indicate that participants held some sort of belief, even if implicit, that there was something hidden and internal to a raccoon, for example, that made it stay a raccoon despite undergoing changes to its superficial proprieties.

More support for psychological essentialism can be said to come from a series of findings that show natural kind categories are subject to what Gelman (2003) calls "boundary intensification" (Diesendruck and Gelman, 1999; Malt, 1990). By this, she means to refer to a sort "line drawn in the sand" phenomenon such that when we make judgments about whether or not something is a category member of a natural kind (but not artifacts), we judge the thing as either being in-category or outside-category but importantly not in-between. In other words, we judge there to be no such thing as kinda-sorta-maybe being a penguin or water or whatever, it just is or it isn't. There is, however, such a thing as kinda-sorta-maybe being a member of an artifact kind like, for example, a chair, and we tend to judge whether or not something is more or less

a category member of CHAIR on the basis of typicality, that is, the extent to which it resembles typical instances of chairs (Diesendruck and Gelman, 1999).[2]

As is the case with most views, they have critics and psychological essentialism is not immune to this. Other studies have begun to call into question the strength that this proposed intuitive folk belief has on our categorization judgments (Malt, 1994; Tobia et al., 2020). Although these studies were primarily focused on the example of water, they seem to suggest that we sometimes have access to more than just an "essentialist concept" (e.g., WATER = H_2O modulo impurities) when offering categorization judgments for natural kinds.[3] This work also seems to be bolstered by more recent hypotheses about the relationship between lexical items, word meanings, and concepts, specifically, that it is not an isomorphic relationship. (Pietroski, 2017, 2018). The studies that I present below add to this line of criticism against psychological essentialism and, ultimately, these data argue that the previous results which had been chalked up as evidence in favor of the view are actually an effect of a certain experimental task setting. In the general discussion, I will present a way of accounting for these data on what I call the "one meaning–potentially many concepts" view (henceforth OMPC). As straightforward as the view may sound, I will also present some theoretical arguments in an effort to bolster OMPC. After all this has been set out, I will conclude the general discussion with some potential directions for future research.

[2] e.g., I am more willing and not hesitant at all to judge my desk chair as being a category member of CHAIR than I am willing to judge those enormous green lawn chairs at Rutgers Gardens as being a category member of chair. As it goes, they're kinda-sorta-maybe chairs.

[3] The reason these studies mainly focused on water was not arbitrary. Putnam (1975) made talk of water fashionable when he devised his famous Twin Earth thought experiment. For those unfamiliar with this, I suggest reading Putnam (pp. 223 - 227) or going to the appendix of this work which includes our experimental materials and reading the vignette for water.

A short note on the experiments

Since I anticipate some readers will have skipped the primer section, I want to begin this brief note by calling attention to an important methodological inference that gets discussed there. Without getting into too many details, in theory, concepts underwrite higher cognitive competences and so, on the basis of that, we can infer that if we study the outputs of higher cognitive competences in certain experimental task settings, then those results will decidedly tell us something about the conceptual content associated with the concepts serving as inputs. There are a number of ways one could go about studying the "outputs," so to speak, of higher cognitive competences (Machery, 2011), but one fashionable way to do so is by eliciting judgments from participants in response to various kinds of vignettes. This is the approach taken in the experiments below.

For purposes of the present work, I am narrowly interested in categorization and as a result, the present studies only examine participants' categorization judgments. As mentioned in the introduction, psychological essentialism claims that certain concepts, let's call them *essentialist concepts*, are constructed against the background of an intuitive folk belief that members of those corresponding categories possess underlying essences. According to this belief, these underlying essences are supposed to determine the category identity (with respect to an essentialist concept) of any given thing we might wish to categorize, such that whether or not that thing possesses the given essence determines whether or not it is a member of the given category.

That might be confusing on a first gloss, but the predictions on behalf on psychological essentialism are rather straightforward: if one is offering a categorization judgment of something with respect to an essentialist concept, then, if that thing under consideration has the corresponding essence property, it should be judged as category

member. On the other hand, if that thing does not have the corresponding essence property, then it should not be judged a category member. Within the existing work cited earlier, there have been a number of findings suggesting, as Gelman (2003) puts it, "[that] although essencelike construals can be found broadly across domains, they are markedly more consistent and compelling for categories of natural entities" (p. 152). In other words, our concepts for natural kinds (but not artifact kinds) are structured as essentialist concepts, resulting in, again, as Gelman (2003) puts it, "evidence for essentialism [that] is (a) unambiguous for natural kinds and many social kinds, and (b) decidedly more mixed for artifacts" (p. 138). The aim of the present studies, then, is to determine whether people's natural kind categorization judgments are as uniform and essentialist as psychological essentialism would predict them to be.

To test these predictions, the test item vignettes for each of the experiments described entities that lacked the essence-like property(ies) associated with a particular category but nevertheless had all of the superficial or perceptual properties (e.g. the Twin Earth liquid). I refer to this sort of vignette structure as DE SS for different essential–same superficial. Alternatively, in some of the control items, participants received vignettes that were DE DS (different essential–differential superficial) and SE SS (same essential–same superficial). What's more, the test item vignettes for each experiment followed the same within-subject design. For any given test item vignette, it will first be in one of three categories: animate natural kind, inanimate natural kind, or artifact kind. Following this, there are two further distinctions among the vignettes within each category. For animate and inanimate natural kinds, this is a distinction at the level of "vignette style, " more specifically, that between "rooban style" and "twin earth style." I'll elaborate more on this difference in a moment, but it is important to note the vignettes in the artifact kind category lack this distinction. Instead, those vi-

gnettes feature a distinction at the level of the candidate essential property described by the vignette, more specifically, a distinction between the essential property being intention or function. So understood, this culminates in a vignette design structure with three factors (animate natural kind, inanimate natural kind, artifact kind) at one level and within each of those factors, another level with two factors (style: rooban, twin earth; essence-artifacts only: intention, function).

Before wrapping up this short note, I want to return to the distinction at the level of vignette style that I just mentioned. The motivation behind issuing this distinction was to incorporate findings from Ahn et al. (2000) that showed we sometimes treat causal properties as "essence-like, " that is, as properties that determine category membership, even if they aren't necessarily described as being hidden or internal. For example, in one item set, participants learned of the following novel category: "animals called 'roobans' tend to eat fruits, have sticky feet, and build nests on trees," and then were provided with additional supplementary information that eating the fruits happens to cause roobans to have sticky feet which, in turn, happens to cause them to be able to climb trees and build nests. When participants were compared to others who did not receive the supplementary information about the causal background of the above features, they appeared less likely to categorize a novel animal as a rooban if it did not possess the initial causal property (eating fruit) despite being identical in all other ways. The so called "rooban style" vignettes, then, are modeled after Ahn et al. (2000)'s explicit mentioning of causal relations among features, whereas the "twin earth style" vignettes lack this (see appendix for vignettes). With all of that explanation out of the way, we can now wrap up this note and move on to the experiments.

Experiment 1: Forced-choice categorization

According to psychological essentialism, when participants are asked to provide categorization judgments for potential category members in DE SS contexts, they should uniformly respond "is not a category member of x" for natural kinds but not artifact kinds. Rather, for artifact kinds, psychological essentialism would predict that their choices turn out to be more mixed. That much should already be clear from the discussion so far. In this experiment, I use a forced-choice categorization judgment task to test these predictions.

Method

PARTICPANTS. Forty undergraduate linguistics students from Rutgers University-New Brunswick. Participants received extra-credit towards a linguistics course upon completion of the experiment. Two participants were excluded for not meeting the minimum time requirement of 10 minutes, one participant was excluded for taking too long, and another was excluded for not finishing the task.

MATERIALS AND PROCEDURE. Participants were presented with vignettes following the experimental paradigm outlined above as part of a forced-choice categorization judgment task. Each participant was presented with twenty vignettes in a pseudo-randomized order and in their response to each vignette, they were instructed to select the statement they agreed with the most among four possible choices (always following the same order):

(1) This x is [BLANK]

(2) This x is not [BLANK]

(3) This x is both [BLANK] and another kind of x

(4) This x is neither [BLANK] nor another kind of x
(e.g., x = 'liquid' and [BLANK] = 'water')

Twelve of the vignettes were test items which, in accordance with the experimental paradigm, consisted of four animate natural kind items (two rooban style, two twin earth style), four inanimate natural kind items (two rooban style, two twin earth style), and four artifact kind items (two intention, two function). The remaining eight vignettes were control items in DE DS and SE SS contexts (one of each for animate natural kinds, inanimate natural kinds, function artifacts, and intention artifacts).

Results and discussion

As will be the case for all the experiments presented in this work, the data analysis is ongoing. As a matter of fact, I haven't been able to run any statistical analyses on these data or compile them into any visualizations. Right now, all I can present is the cleaned data I gathered in the form of a table:

		Natural kinds				**Artifact**	
		Animate		**Inanimate**		**Intention**	**Function**
		Roobans	**Twin earth**	**Roobans**	**Twin earth**	**n/a**	**n/a**
1	is an X	**38.75%**	20.00%	**35.00%**	18.75%	**37.50%**	21.25%
2	is both an X and a Y	5.00%	20.00%	8.75%	10.00%	**36.25%**	**40.00%**
3	is neither an X nor a Y	2.50%	2.50%	1.25%	3.75%	3.75%	5.00%
4	is not an X	**53.75%**	**57.50%**	**55.00%**	**67.50%**	**22.50%**	**33.75%**
		100.00%	100.00%	100.00%	100.00%	100.00%	100.00%

Test items for Experiment 1

		Natural kinds				Artifact			
		Animate		Inanimate		Intention		Function	
		SE_SS	**DE_DS**	**SE_SS**	**DE_DS**	**SE_SS**	**DE_DS**	**SE_SS**	**DE_DS**
1	is an X	**90.00%**	2.50%	**87.50%**	2.50%	**77.50%**	7.50%	**72.50%**	5.00%
2	is both an X and a Y	5.00%	15.00%	2.50%	5.00%	20.00%	10.00%	15.00%	20.00%
3	is neither an X nor a Y	2.50%	2.50%	0.00%	5.00%	0.00%	5.00%	2.50%	2.50%
4	is not an X	2.50%	**80.00%**	10.00%	**87.50%**	2.50%	**77.50%**	10.00%	**72.50%**
		100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%

Control items for Experiment 1

Notably, the majority of participant responses for natural kinds, both animate and inanimate, was "is not an x, " which was perfectly in line with the predictions from psychological essentialism. Likewise, participants' responses for artifacts were comparatively more variable than their responses to natural kinds, which was also in line with the predictions from psychological essentialism.

On a first gloss, then, the results from Experiment 1 appear to replicate the findings from previous work that have been put forward as evidence in favor of psychological essentialism. However, on a second gloss, we should notice that although the majority of participant responses for natural kinds were essentialist, they were not as uniform as we might have expected. In some case, the majority of participant responses was only a little bit more than half and so, on that note, Experiment 2 aims to offer more scrutiny towards these early results.

Experiment 2: Truth value judgments

The above observations only shed light on participants' categorization judgments in a very specifically designed forced-choice task where they are only permitted to make a judgment in the form of one choice. To obtain additional evidence about the strength

of this purported intuitive folk belief on our categorization judgments, Experiment 2 sought to examine participants' judgments in a less restricted experimental task setting. Since participants are still offering categorization judgments and the only change being made between Experiment 1 and Experiment 2 is a change in the task, the predictions from psychological essentialism should remain the same as before.

Method

PARTICPANTS. Sixty-seven undergraduate linguistics students from Rutgers University-New Brunswick. Participants received extra-credit towards a linguistics course upon completion of the experiment. Four participants were excluded for failure to finish the study.

MATERIALS AND PROCEDURE. Participants provided categorization judgments as part of a two-block truth-value judgment task. The first block consisted entirely of test item vignettes in the DE SS context (the same twelve vignettes from Experiment 1 following the same design structure and pseudo-randomization procedure). After reading each vignette, participants were presented with four statements and for each statement, they were instructed to indicate whether they thought it was true or false. The four statements were the same statements that were originally presented as choices in the forced-choice task from Experiment 1.

The second block consisted of some control item vignettes in the DE DS and SE SS contexts and some test item vignettes in the DE SS and SE DS contexts. SE DS was not examined in Experiment 1, so Experiment 2 purposefully implemented vignettes to test this context. The vignettes in this second block featured stories about artifacts undergoing transformations similar to the vignettes used in Keil (1989)'s ex-

periment presented earlier. These vignettes were constructed in a four different item types (teakettle, birdhouse; garbage can, chair; PVC pipe, flute; cards, wipes) x four different contexts (DE DS; DE SS; SE SS; SE DS) design. The vignettes were randomized using a latin square, such that participants received a total of four vignettes (one for each context) and the item types did not overlap.

Results and discussion

I'm mentioning this again: data analysis is still ongoing. For the time being, I can only present cleaned data in the form of tables.

	Natural kinds				Artifact	
	Animate		Inanimate		Intention	Function
	Roobans	Twin Earth	Roobans	Twin Earth	n/a	n/a
is an X	**72.6%**	33.9%	39.5%	18.5%	**61.3%**	**53.2%**
is not an X	32.3%	**62.9%**	**56.5%**	**79.8%**	**47.6%**	**52.4%**
both an X and a Y	33.9%	37.9%	21.8%	25.0%	**54.0%**	**41.9%**
neither an X nor a Y	2.4%	8.1%	11.3%	8.1%	4.0%	8.1%

% of responses for true - block one, Exp. 2

	Transformation			
	DE_DS	SE_DS	DE_SS	SE_SS
is an X	19.4%	**80.6%**	38.7%	**93.5%**
is not an X	16.1%	22.6%	**79.0%**	30.6%
both an X and a Y	8.1%	35.5%	35.5%	45.2%
neither an X nor a Y	**71.0%**	19.4%	12.9%	0.0%

% of responses for true - block two, Exp. 2

On a first gloss of these results, we notice that a majority of "true" responses for natural kinds are for "is not an X" as psychological essentialism predicts. Likewise, on our first gloss, we notice responses are more mixed for artifacts, which is also what psychological essentialism predicts. On a second gloss, however, we notice that par-

ticipants also offer a high number of "true" responses for other statements as well. Taking those results into consideration, then, these data appear to weigh against the predictions from psychological essentialism. More specifically, they seem to indicate that despite largely responding "true" to the essentialist option, participants are not also responding "false" to all the other options. Quite the opposite.

Experiment 3: Faultless Disagreement

In the two experiments reported so far, participants were asked to provide their own categorization judgments in response to various kinds of vignettes. The main observation so far has been that participants' judgments appear in line with the predictions from psychological essentialism when they are in a particular experimental task setting, however, when the experimental task setting allows for less restricted judgments, then these judgments no longer appear in line with those predictions. This experiment seeks to ask a follow-up question: what would happen if participants were instead asked to judge the categorization judgments of others?

To test this, I presented participants with vignettes as part of a faultless disagreement task. Following each vignette in this task, participants were shown categorization judgments from two different speakers. In the prompts, the speakers were referred to as Speaker A and Speaker B, but for purposes of elucidating the experimental design, it might be easier to refer to Speaker A as Superfical Sam and Speaker B as Essentialist Elaine. In response to every vignette, Elaine always categorizes on the basis of the essential property(ies) and Sam always categorizes on the basis of the superficial properties. So, in other words, in a DE SS context, Elaine always responds "is not an X" and Sam always responds "is an X." According to psychological essentialism, one

would expect people to always judge Elaine as being right and Sam as always being wrong.

Method

PARTICPANTS. Forty-one undergraduate linguistics students from Rutgers University-New Brunswick. Participants received extra-credit towards a linguistics course upon completion of the experiment. Six participants were excluded for failure to meet the minimum time requirement of 15 minutes.

MATERIALS AND PROCEDURE. Two block forced-choice faultless disagreement task with adjusted versions of the already created vignettes for animate natural kinds, inanimate natural kinds, and artifacts. The structure of the blocks is identical to Exp 2. Block one features test items and faultless disagreement control items (no transformation block), and block two features transformation scenarios, which here act as controls for essentialist reasoning (that is, people's responses in DE DS and SE SS contexts should reflect how they responded in Experiment 1). The faultless disagreement controls consisted in four vignette types, including one in which Speaker A was undoubtedly right, one in which Speaker B was undoubtedly right, and two in which faultless disagreement should be allowed (relative gradable adjectives and predicates of personal taste). After being presented with the vignettes, participants are instructed to choose one of four mutually exclusive answers

(1) A is RIGHT and B is WRONG.
(2) A is WRONG and B is RIGHT.
(3) There is a sense in which both A and B are RIGHT.
(4) Neither A nor B is right.

Results and discussion

As before, data analysis is still ongoing. For the time being, I can only present cleaned data in the form of tables.

	Natural kinds				Artifact	
	Animate		Inanimate		Intention	Function
	Roobans	Twin Earth	Roobans	Twin_earth	n/a	n/a
1. A is right and B is wrong.	28%	28%	19%	6%	**44%**	14%
2. A is wrong and B is right.	33%	**44%**	**58%**	**72%**	17%	33%
3. Both are right.	31%	25%	19%	17%	39%	**53%**
4. Neither A nor B is right.	8%	3%	3%	6%	0%	0%

Test items for Experiment 3

	Transformation			
	DE_DS	SE_DS	DE_SS	SE_SS
1	5.6%	**27.8%**	16.7%	**66.7%**
2	**83.3%**	**44.4%**	**77.8%**	11.1%
3	5.6%	**27.8%**	5.6%	16.7%
4	5.6%	0.0%	0.0%	5.6%
	100.0%	100.0%	100.0%	100.0%
			Control	
	A_right	B_right	RelGA_FD_ok	PPT_FD_ok
1	**86.1%**	2.8%	2.8%	11.1%
2	8.3%	**91.7%**	0.0%	5.6%
3	5.6%	5.6%	**94.4%**	**75.0%**
4	0.0%	0.0%	2.8%	8.3%
	100.0%	100.0%	100.0%	100.0%

Control items for Experiment 3

Immediately, we notice that a majority of participant responses are for the option that Elaine is right and Sam is wrong. In other words, people seem to be agreeing with the essentialist judgment. However, on a closer look, there is considerable variability in how participants respond across conditions for animate and inanimate natural kind

items and rooban and twin earth items. In fact, participants seem to be pretty split between options 1 - 3 for the rooban animate natural kind items.

Even within the condition that received the highest percentage of responses for option two, there appears to be considerable variability in participants' responses when we compare the different items. In particular, when examining the "water" and "helium" vignettes, we see almost one-third of participants selecting option three in response to the "water" item but not in response to the "helium" item.

	Inanimate Twin Earth	
	Helium	Water
1	5.6%	5.6%
2	**77.8%**	**66.7%**
3	5.6%	27.8%
4	11.1%	0.0%
	100.0%	100.0%

General discussion

The rest of this work will be dedicated to pulling out the takeaways from the above experiments. Accordingly, this discussion will be split this up into three subsections. In first subsection, I will discuss general conclusions about psychological essentialism that I believe are revealed by these experiments. The way I see it, the above data are completely at odds with psychological essentialism as it has been traditionally formulated, however, I suggest we might not have to throw out the baby with the bathwater. Following this, the second subsection will address how these results bear on deeper questions concerning mental architecture, conceptual content, and conceptual structure. In that section, I will argue that the OMPC view can account for the results from

above experiments and offer an alternative way of formulating psychological essentialism such that it is consistent with these results. Finally, the third subsection will highlight some potential directions for future research.

Conclusions about psychological essentialism

This paper began by calling attention to previous work which has been said to demonstrate that some of our concepts, particularly, our natural kind concepts, are constructed against the background of an intuitive folk belief that the things which we mentally represent with those concepts have essences (Ahn et al., 2000; Diesendruck and Gelman, 1999; Gelman, 2003; Gelman and Wellman, 1991; Keil, 1989; Malt, 1990; Medin and Ortony, 1989). This previous work was also said to demonstrate that artifact kind concepts are different in this regard. As the story goes, then, since conceptual content for natural kinds (but not artifact kinds) is constituted by this intuitive folk belief and since, according to our *pro tem* theory of concepts, concepts underwrite higher cognitive competences like categorization, it follows, we should show different patterns in our categorization judgments of natural kinds and artifact kinds. More specifically, our categorization judgments for natural kinds should always reflect a change in category identity in DE SS contexts, whereas with artifacts, our judgments should be more mixed. In a replication of the previous work and a demonstration of precisely this line of reasoning, a majority of participant responses in Experiment 1 were essentialist for natural kinds but not artifact kinds.

Here's the problem though: most, if not all, of the previous work (including Experiment 1) was conducted as a forced-choice task in which participants could only choose a "simple" categorization judgment. To elaborate on what I mean by "simple," I

want to point out that in Experiment 1, participants were only allowed to choose one of the following four statements in response to a vignette.

(1) This x is [BLANK]
(2) This x is not [BLANK]
(3) This x is both [BLANK] and another kind of x
(4) This x is neither [BLANK] nor another kind of x
(e.g., x = 'liquid' and [BLANK] = 'water')

Choosing only one of these statements is simple as opposed to being complex because one can imagine having a categorization judgment that incorporates statements (1) and (3) as a conjunction. Endorsing such a categorization judgment is very different from endorsing only (1) or (3) in isolation. For example, if we consider these judgments in the context of the Twin Earth thought experiment, someone who judges "(1) & (3)" might think there is a sense in which the twin earth liquid is water and that there is a sense in which it is not water but, for what its worth, lean more on the side of thinking its water than thinking its not water. On the other hand, someone who judges (1) or (3) in isolation judges one or the other, that is to say, they don't endorse some judgment found in the middle. However, as the findings from Experiment 2 and Experiment 3 demonstrate, when this issue is taken into consideration and participants are allowed to offer more complex categorization judgments, their responses are less uniform and, in some cases, no longer essentialist.

We might consider these results as hinting at a rule of thumb for experimental design in behavioral science, that is, just because a participant selects one particular response in a forced-choice task does not necessarily mean they are responding "no" to all the other choices. Like we just reviewed, if particiants are given an opportunity,

they might select responses (1) and (3) from above and that would be perfectly reasonable not to mention an importantly different kind of response than selecting either (1) or (3) in isolation would be. We also might consider the results presented here as suggesting that psychological essentialism does not offer a complete picture of the role that essentialist beliefs play in (i) our concepts for certain natural kind categories and (ii) our categorization judgments. As devastating as that may seem for proponents of psychological essentialism, one does not need to throw the baby out with the bathwater. In fact, I think Bloom (2007) largely anticipates these problems and offers a view of what he calls "hybrid concepts" as a sort of rejoinder.

According to this view, some categories are not mentally represented as just being a natural kind or just being an artifact kind but rather as a natural kind–artifact kind hybrid. For example, Bloom seems to suggest that the conceptual content associated with a concept like WATER is a conjunction of the conceptual contents associated with WATER:SCIENTIFIC and WATER:FUNCTION concepts.[4] On this way of viewing things, then, one could maintain that essentialist beliefs play some sort of role in our categorization judgments since, in theory, they play a role in the conceptual content associated with a concept like WATER:SCIENTIFIC, but also maintain that because some categories are represented as hybrid concepts with hybrid content, there can sometimes be variability in how we categorize to the point that we might not always reflect essentialist judgments for what we would have traditionally considered to be natural kinds. How exactly a categorization procedure using hybrid concepts results in different judgments given that, although they are hybrid, they are still monadic, could be cast off as a separate

[4] WATER:SCIENTIFIC refers to the essentialist concept of water being H_2O modulo impurities and WATER:FUNCTION refers to a functional concept of water. On a functional concept of water, it is just a liquid possesses the functional role of water, namely, something that comes out of the sink faucet, quenches thirst, powers hydroelectric machines, and so on.

issue. All that a person offering this sort of rejoinder needs to show, at least in this setting, is that there is room in psychological essentialism to account for the above data.

Now, to be sure, I would consider a rejoinder of this sort a step in the right direction for reasons that will become clear in the next subsection, but I believe it faces some pitfalls. For one thing, a view of hybrid concepts faces what I'll call the "invalidity of inferences problem." Basically, the issue is this: when something is categorized as WATER:FUNCTIONAL it does not license a valid inference that it can also be categorized as WATER:SCIENTIFIC, however, according to a theory of hybrid concepts that claims they have hybrid content, this should be a valid inference. The Twin Earth thought experiment is a perfect case in point. Just because we are willing to categorize the twin earth liquid as WATER:FUNCTIONAL does not mean we can also infer that it can be categorized as WATER:SCIENTIFIC. In fact, according to psychological essentialism, it cannot be categorized as WATER:SCIENTIFIC because it is not chemically composed of H_2O and therefore lacks the essential property. But, under a theory of hybrid concepts, if we categorize the twin earth liquid as WATER:HYBRID, we appear committed to categorizing it as both WATER:FUNCTIONAL and WATER:SCIENTIFIC.

To repeat, although I don't think a theory of hybrid concepts is plausible, we still don't need to throw the baby out with the bathwater when it comes to psychological essentialism. In the next subsection, I will present an alternative theory to hybrid concepts, namely OMPC, that could account for the above data.

One word meaning – potentially many concepts

I mentioned that I consider the theory of hybrid concepts to be a step in the right direction and, to begin this subsection, I wish to elaborate on why I think that. At this point it should be fairly clear why the above data are at odds with psychological essentialism, but in case it's not, I'll quickly summarize the point: when people are given the opportunity to offer complex categorization judgments in DE SS contexts, their responses are no longer as uniform or essentialist as psychological essentialism would predict them to be. A possible way of accounting for these findings might be to argue that when we are offering categorization judgments for natural kinds, we are not merely using "essentialist concepts" or, in other words, concepts that are constructed against the background of an intuitive folk belief in essences. In that case, we will want to have a theory about how this is possible. The theory of hybrid concepts from Bloom (2007) was on attempt at this, but the for the reasons I just presented in the last subsection, it is not plausible. Thus, in that way it was a step in the right direction, it's just that not every step we take, even if in the right direction, gets us to where we want to be.

There is, I think, a traditionally uncontested assumption being made here which is blocking us from our path to where we want to be. More specfically, this assumption is that the relationship between lexical items, word meanings, and concepts is a one-to-one-to-one correspondence such that any given word only has one concept. Putting aside the issue of why this assumption has been traditionally uncontested, when we move away from thinking of the above relationship as a one-to-one-to-one correspondence, we notice that accounting for the findings presented in this paper becomes quite simple. According to OMPC, lexical items and word meanings have a one-to-one

correspondence, but word meanings are thought of in a nontraditional and perhaps unusual way. For propoents of OMPC, word meanings are just instructions on where to locate a particular concept in our mental lexicon. For our present purposes, these get called "fetch@address-in-lexicon" instructions and at any given address, there are potentially many concepts. At any rate, the point to make here is that if a single word can be connected to potentially many concepts, then one of those concepts might be an essentialist concept, another one of those concepts might a functional concept, and so on. When we make categorization judgments, then, we might be using any one of those concepts in our categorization procedure – it doesn't always need to be the essentialist concept. If we wish to apply this theory to the data presented in Experiments 1 - 3, we might do so by saying the mixed responses we find are explained by participants sometimes categorizing on the basis of a non-essentialist concept.

In case it isn't clear, I want to point out that supporters of the OMPC view still hold onto a version of psychological essentialism, albeit, in a much weaker form. On OMPC, one still maintains that some concepts are constructed against the background of an intuitive folk belief in essences and, as a matter of fact, one ought to still maintain this in light of the previous work that was cited as evidence in favor of psychological essentialism. In other words, there is no getting away from the idea that some concepts are essentialist, however, we do want to get away from the idea that all or even most natural kind concepts (and no artifact kind concepts) are essentialist. With OMPC in mind, words are not inextricably linked to only one concept and, as a result, we can use different concepts in our categorization procedures.

Besides being able to account for the above data, I think there are interesting theoretical arguments that could be made in favor of OMPC (Pietroski, 2018). I will present one argument here, but in order to set this up, we shall want to first consider the fol-

lowing case involving co-predication:

> *Imagine a logician who can't remember whether it's France or Switzerland that is hexagonal, but is sure that of those two, the hexagonal country is the one that is also a republic. Suppose he reasons as follows: if France is hexagonal, then it is a republic; so if France is hexagonal, then it is both hexagonal and a republic; so if nothing is both hexagonal and a republic, then France isn't hexagonal and Switzerland is. But likewise, if Switzerland is hexagonal, it's a republic; so if Switzerland is hexagonal, it's both hexagonal and a republic; however, if nothing is both hexagonal and a republic, Switzerland isn't hexagonal and France is. Given that, it must be false that nothing is both hexagonal and a republic. So, in other words, something—France or Switzerland—must both be hexagonal and a republic.*[5]

This case is a modification of a fairly classic example used to illustrate polysemy. For the logician in the story, the confusion arises from making the following logical move from (i) to (ii).

(i) $Hf \supset Rf$

(ii) $Hf \supset (Hf \,\&\, Rf)$

Although the move might initially seem innocuous when it's presented in the story, we know independently that the sort of things which are republics are not the sort of things which can be hexagonal. Strictly speaking, governments cannot be hexagonal and objects shaped like hexagons are not republics. That is to say, the existence claim $(\exists f)(Hf \,\&\, Rf)$ is false, there is no unitary thing, f, that has both the property of being

[5] This is an unpublished example from Paul Pietroski.

hexagonal and a republic. But, if that's true, then how is the logician supposed to figure out whether it is France or Switzerland that is hexagonal provided he knows the hexagonal country is also the one that is a republic?

Well, it turns out to be possible for the logician provided he understands that the lexical item 'France' is polysemous. 'France' denotes (i) a particular geographical region in Western Europe and (ii) a particular government. So, that being the case, in the geographical sense, France can be hexagonal and in the sense of government, France could be a republic and all of this could be true without the logician needing to be committed to the claim that there exists a unitary thing which is both hexagonal and a republic. That said, here's the main line of the argument: if lexical items, word meanings, and concepts had an isomorphic relationship, then it would be impossible for us to make sense out of why statements like "France is both hexagonal and a republic" are felicitous. To show this, consider what the result would be from thinking there was an isomorphic relationship and that, in fact, the above story was an instance of homophony instead of polysemy.[6] That would imply that what is occurring with 'France' in the above story is an instance of there being two different words with the same pronounciation, but if that's so, then there's no way of accounting for how statements like "France is both hexagonal and a republic" are felicitous. Since the common example of homophony involves 'bank' the river edge, and 'bank' the financial institution, we should consider how a statement like "Bank is both near a river and a financial institution" is seemingly infelicitous despite being a similar construction involving copredication. Clearly what is going on with 'France' is not the same sort of phenomenon as what is going on with 'bank' and 'bank'. To imitate Putnam (1975), cut the pie any

[6]Homophony is the phonomenon of two different words (and correspondingly, pace OMPC, two different concepts that are associated with each word) having the same pronouncation like, for example, 'Bank' referring to the edges of a river and 'Bank' referring to a financial institution.

way you like, the picture must be one word meaning – potentially many concepts.

Future directions

Notwithstanding that I hold the above theoretical argument to weigh decisively in favor of OMPC, I must add that Experiments 1 - 3 are not able to adjudicate between OMPC and something else. Although a theory may be able to account for a set of findings, it is quite a different thing for a set of findings to actually adjudicate between competing theories. As it happens, this is just the position we are currently in when it comes to examining OMPC. We have some exciting and overwhelming theoretical support but, alas, no real empirical support in its favor. On that note, then, this where I see a promising direction for future research.

The first step to testing a theory is getting clear about what its predictions are. With respect to OMPC, I'll attempt to outline the predictions here, but I'll want to do so against the background of the Twin Earth thought experiment. In the context of this case, OMPC claims that the lexical item 'water' has a single fetch@address-in-lexicon instruction attached to it in the form of a meaning and, at this address, there are at least two concepts, namely, WATER:SCIENTIFIC and WATER:FUNCTIONAL. Given this, it should be possible for one to provide categorization judgments using both concepts in the Twin Earth thought experiment, that is, one should be able to categorize the twin earth liquid as 'water' under the WATER:FUNCTIONAL concept but not under the WATER:SCIENTIFIC concept. There is a way of interpreting the results from Experiment 3 as already giving some credence to this thought. To recall, this experiment used a faultless disagreement paradigm in which Superficial Sam was described as offering a categorization judgment using the WATER:FUNCTIONAL concept and Essentialist

Elaine was described as offering a different categorization judgment using the WATER:SCIENTIFIC concept, but regardless of the differences in their judgments, roughly one-third of participants still responded that that both judgments were correct.

There is, however, one major reason why Experiment 3 can only be interpreted as giving some credence to OMPC but not as providing any decisive evidence for or against it. This is because the experimental design as it currently stands does not ensure that participants are thinking of Sam and Elaine's judgments as involving two different concepts. They might be, but we just can't be sure of it. I don't think we are in a totally hopeless position though, the way I see it Experiment 3 can be modified to control for this problem. Here's how I envision the modification: after Sam and Elaine present their judgments in light of the case information they will present justifications for their judgments. For example, after Sam judges the twin earth liquid to be water he will offer a justification for that claim by saying something to the effect of the following: "the twin earth liquid comes out of sink faucets, is sourced from rivers and lakes, powers hydroelectric machines, quenches thirst, etc., just like water on earth does; therefore, it is water." *Mutatis mutandis* for Elaine. By offering justifications for each speaker's categorization judgment, it should fairly evident to participants that Sam and Elaine (literally) have different conceptions of water in mind. If after making this adjustment to the experiment it is discovered that a significant percentage of participants respond that both Sam and Elaine's judgments are correct, then, I think, this will be decisive evidence in favor of OMPC.

Acknowledgments

This research was completed over the course of the 2020-2021 academic year in which the entire global community was battling the coronavirus pandemic. Every university was closed and everything was virtual for what looked like the foreseeable future. At times, too, the pandemic seemed to distract us from the fact that 2020 featured a massively divisive election year in U.S. , so much so that back in January a coup of insurrectionists actually stormed the capital building. I don't mean it lightly when I say that on some days I wouldn't have been surprised to wake up to a news notification that read "breaking: Earth stops spinning on its axis." In a year like that, one would only expect that productivity is flung to the wind and according to philosophy and cog sci Twitter, most academics experienced this. All of which is in an effort to express how indebted I am to my co-supervisors, Kristen Syrett and Paul Pietroski, for their unwavering dedication and professionalism that kept me on track. As I remember Paul likes to say, "the trains must run on time, " and it is because of them that the above work got to its station.

References

Ahn, W.-k., Kalish, C., Gelman, S. A., Medin, D. L., Luhmann, C., Atran, S., Coley, J. D., and Shafto, P. (2001). Why essences are essential in the psychology of concepts. *Cognition*, 82(1):59–69.

Ahn, W.-k., Kim, N. S., Lassaline, M. E., and Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive psychology*, 41(4):361–416.

Bloom, P. (2007). *Creations of the Mind: Theories of Artifacts and Their Representations*, chapter Water as an Artifact Kind. Oxford University Press.

Chomsky, N. (1995). Language and nature. *Mind*, 104:1+.

Clark, E. (1983). Meanings and concepts. *Handbook of child psychology ; Cognitive Development*, 3.

Diesendruck, G. and Gelman, S. A. (1999). Domain differences in absolute judgments of category membership: Evidence for an essentialist account of categorization. *Psychonomic bulletin I& review*, 6(2):338–346.

Fodor, J. (1998). *Concepts: where cognitive science went wrong*. Oxford University Press, Oxford.

Gelman, S. (2003). *The essential child : origins of essentialism in everyday thought*. Oxford University Press, Oxford.

Gelman, S. and Wellman, H. (1991). Insides and essences: Early understandings of the non-obvious. *Cognition*, 38(3):213–244.

Gelman, S. A. (2013). Artifacts and essentialism. *Review of philosophy and psychology*, 4(3):449–463.

Keil, F. (1989). *Concepts, kinds, and cognitive development*. MIT Press, Cambridge, MA.

Komatsu, L. K. (1992). Recent views of conceptual structure. *Psychological bulletin*, 112(3):500–526.

Machery, E. (2011). Concepts: a tutorial. In Belohlavek, R. and Klir, G. J., editors, *Concepts and Fuzzy Logic*. The MIT Press.

Malt, B. C. (1990). Features and beliefs in the mental representation of categories. *Journal of memory and language*, 29(3):289–315.

Malt, B. C. (1994). Water is not h2o. *Cognitive Psychology*, 27(1):41 – 70.

Malt, B. C. and Sloman, S. A. (2007). *Creations of the Mind: Theories of Artifacts and Their Representations*, chapter Artifact categorization: The good, the bad, and the ugly. Oxford University Press, Oxford.

McCloskey, M. and Glucksberg, S. (1978). Natural categories: Well defined or fuzzy sets? *Memory I& Cognition*, 6:462–472.

Medin, D. and Ortony, A. (1989). *Similarity and analogical reasoning*, chapter Psychological essentialism, pages 179–195. Cambridge University Press.

Pietroski, P. M. (2017). *The Cambridge Companion to Chomsky*, chapter Semantic Internalism. Cambridge University Press.

Pietroski, P. M. (2018). *Conjoining meanings: Semantics without truth values.* Oxford University Press.

Putnam, H. (1975). The meaning of meaning. *Minnesota Studies in the Philosophy of Science: Language, Mind, and Knowledge*, 7:131–193.

Rosch, E. and Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7:573–605.

Shea, N. (2018). *Representation in Cognitive Science.* Oxford University Press, Oxford.

Tobia, K. P., Newman, G. E., and Knobe, J. (2020). Water is and is not h20. *Mind I& language*, 35(2):183–208.

Wittgenstein, L. (1953). *Philosophical Investigations.* Macmillian, New York.

Appendix to all experiments

1. No Transformation: Same perceptual features, different essential properties

1.1. Natural kind: Animate

1.1.1. Twin Earth Style

1.1.1.1. Reptile Imagine a newly-discovered animal that, upon inspection, resembles reptiles in all of its outward properties. It has dry, scaly skin, a backbone, a long tongue, eats insects, and is found living in warm, sunny habitats. It also lays eggs instead of live young, and these eggs are rubbery and initially white. However, upon further examination, scientists discover that its circulatory system is different from all known reptiles, since it lacks a divided atrium, and its blood typically has more hemoglobin.

This animal is a reptile. This animal is not a reptile. This animal is both a reptile and another kind of animal. This animal is neither a reptile nor another kind of animal.

1.1.1.2. Cow Imagine an animal that resembles a cow in all important outward characteristics. It is moderately large, has hair, four legs, and a tail, emits a sound like a 'moo', and gives milk. Some people even like to attach a sort of bell around its neck that makes a "clanky" sound when this animal moves. However, this animal lacks a ruminant stomach with four compartments (like a cow has), which affects its diet.

This animal is a cow. This animal is not a cow. This animal is both a cow and another kind of animal. This animal is neither a cow nor another kind of animal.

1.1.2. Roobans Style

1.1.2.1. Boskets Biologists have discovered an animal living in the jungle that they call "boskets." Boskets have very poor eyesight but have a keen sense of smell, have small pupils but large nostrils at the end of a long thin nose, and eat fragrant flowers that grow in the jungle. Biologists have argued that because of the low light under the tree canopy in the jungle and the prevalence of these nutrient-rich flowers in their habitat, over time, boskets have developed these physical features for survival. Imagine that biologists discover an animal that has these same physical characteristics, but is native to the pine barrens of NJ.

This animal is a bosket. This animal is not a bosket. This animal is both a bosket and another kind of animal. This animal is neither a bosket nor another kind of animal.

1.1.2.2. Modis Animals called "modis" live in ponds. They eat the grass that grows in the water near the shore, have green colored skin that allows them to blend in with the surroundings, and lay eggs in the brush, safe from the fish in the water that would eat them. After hatching, the young make their way to the water. Biologists have deduced that, over time, the properties of the grass are what led to the green coloration of the skin and, ultimately, their green color endowed them with the camouflage they needed to venture into the brush, allowing them to shield themselves and their future young from predators. Suppose biologists discover an animal that has the following characteristics: it's native to riverbeds, eats tiny insects, has green skin, and lays eggs in the brush on the river bank. The young make their way to the river upon hatching.

This animal is a modi. This animal is not a modi. This animal is both a modi and another kind of animal. This animal is neither a modi nor another kind of animal.

1.2. Natural kind: Inanimate

1.2.1. Twin Earth Style

1.2.1.1. Helium Imagine a gas that is exactly like helium (He) in its observable properties. It is colorless, odorless, tasteless, non-toxic, inert, lighter than air (and so makes balloons float), and has an extremely low boiling point. When you suck on it, it transforms your voice to a high pitch. At the same time, this gas appears to be different in fundamental properties. When chemists examine it, they find that its atomic mass is not 4.002602 amu, and they are uncertain about whether there are 2 protons in the atom's nucleus.

This gas is helium. This gas is not helium. This gas is both helium and another kind of gas. This gas is neither helium nor another kind of gas.

1.2.1.2. Water Scientists have outlined possible scenarios for the event that sometime in the future, other planets are found within a neighboring galaxy. If so, a space expedition would be undertaken to explore these planets. In one scenario, the expedition to one of these planets discovers a liquid on the planet. The liquid is clear and colorless, is drinkable with no strongly detectable taste, and is indistinguishable from the liquid one would find in brooks, ponds, and rivers here in Earth. When it freezes, it becomes a solid. When it heats up, it becomes a gas. At the same time, data analysis reveals that this liquid does not have the exact same molecular structure as the H20 on our planet, but is in all other ways indistinguishable.

This liquid is water. This liquid is not water. This liquid is both water and another kind of liquid. This liquid is neither water nor another kind of liquid.

1.2.2. Roobans Style

1.2.2.1. Kalumin Botanists have discovered a tropical plant with large leaves that bears brightly colored yellow fruit, which a certain kind of tree-dwelling animal finds particularly delicious. When this animal nibbles on the fruit of the plant, their saliva starts a chemical process in the plant that increases its production of a particular vitamin, making the plant extremely useful in medicines. They call this plant a "kalumin." Suppose botanists find another kind of plant that has a highly similar appearance (large leaves, bright yellow fruit, and so forth), which also has a similarly high concentration of this kind of vitamin, and is therefore also used frequently in medicines. However, this other plant grows in a location of the forest that is low on the forest floor in a place that is distinct from the habitat of this particular animal. In fact, it is not clear if any animals eat this fruit.

This plant is a kalumin. This plant is not a kalumin. This plant is both a kalumin and another kind of fruit. This plant is neither a kalumin nor another kind of fruit.

1.2.2.2. Pimwiscus Plants called a "pimwiscus" tend to grow in warm tropical regions. They attract certain kinds of insects into their container-like base to digest them in an acidic fluid, have a waxy, brightly-colored outside, and are found among vibrant flora that attach themselves to the base of trees. Botanists have discovered that the color and texture of the plant is due to the way it digests the insects, and this color in turn allows them to blend in with the vibrant flora, allowing them to attract the insects that fly in for the sugary nectar in these flowers. Suppose botanists find a plant that does not digest insects, has a waxy, colorful exterior, and grows amidst these flora.

This plant is a pimwiscus. This plant is not a pimwiscus. This plant is both a pimwiscus and another kind of plant. This plant is neither a pimwiscus nor another kind of plant.

1.3. Artifact

1.3.1. Intention

1.3.1.1. Pastry chef During a recent pastry making competition in Lille, France, the patissiers (a fancy word for pastry chef) were tasked with crafting a new creation, one that the judges had never seen before and that would wow both their eyes and their taste buds. The chef who won the contest spent considerable time sketching their design for a pastry that has layers of flaky dough shaped into a cross shape, with one cross-section containing a savory, tangy cheese, and the other cross-section containing a tart yet sweet red preserve. The chef's idea was that the lucky taste-tester would bring the two sections together, dip this into a mixture of fig and honey, and take a bite they would never forget. They called this creation a "croix-délice". Around the same time, in a kitchen in Bethesda, MD, a woman named Amelia was talking on the phone to her mom, who happened to call just before lunchtime. As she talked, Amelia distractedly took things out of the fridge and the cupboards to make herself something to eat. She took two long pieces of phyllo dough from leftover Greek food, spread goat cheese on one, and spread the last of some lingonberry preserves on the other. As her mom went on about an upcoming yard sale in the neighborhood, Amelia reached into the fridge, absentmindedly pulled out a jar of fig and honey preserves she got in a gift basket from an office raffle, and dipped the covered strips into it and took a bite as she listened to her mom tell a story about the deer eating up her vegetable garden.

This food preparation is a croix-délice. This food preparation is not a croix-délice. This food preparation is both a croix-délice and another kind of food preparation. This food preparation is neither a croix-délice nor another kind of food preparation.

1.3.1.2. Sculpture An artist was commissioned by a US city to create a sculpture that depicted the year 2020 and how we could emerge from it. The artist first created a model of their sculpture to display for the city council before construction on the larger piece was done. Here's what they did to create the object. First, they took a sturdy beam of wood, sanded it down, polished it, and shaped it into an obelisk form. They dug a hole into an open space into the ground, staked the wooden form into the hole, and filled the hole with cement, so that the obelisk remained sturdy. This symbolized unity and strength, and referenced similar historical monuments that have lasted throughout tumultuous events in the world. The artist then attached a hollow metal object towards the top that was flat at the bottom but had a curved top, symbolizing duality of false claims and flat earthers and science-based reasoning. A red metal flag attached to the right side could be raised and lowered to symbolize the rising and falling of scientific reason, the Republican party, and the sun setting on this year full of anxiety and stress. The container had a door that could be opened and closed so that people could write letters about their experiences and deposit them inside for someone else to read. When an intern walked by, he remarked that the artist's sculpture looked just like a mailbox.

This object is a mailbox. This object is not a mailbox. This object is both a mailbox and another kind of object. This object is neither a mailbox nor another kind of object.

1.3.2. Function

1.3.2.1. Archeological dig Some researchers were recently on an archeological dig, where they found a small light colored object in a grave for a young girl. This object was found near her skull, which made them think that it must have been intended to be worn as ornamental jewelry in her hair. The stone has a round shape with two holes about an inch apart, and a

curved stick passing through the holes. They called this newly found object a "funduletta." A month later, they were on a dig at another site, and were excavating objects in what they determined to be a space where hunters gathered and assembled their kills. In the space where weapons were amassed, amidst the knives, axes, and bows and arrows, the archeologists found a small, round light colored object with two holes about an inch apart, and a curved stick passing through the hole. They determined this object must have been intended to be used to hold the quivers for the bow and arrow.

This object is a funduletta. This object is not a funduletta. This object is both a funduletta and another kind of object. This object is neither a funduletta nor another kind of object.

1.3.2.2. Native Americans Imagine an object that is either made of metal or a hard plastic. On one end, it protrudes narrowly about 10", making it easy to grasp at that end. On the other end, it extends into a flat somewhat rectangular surface about 3" x 4", which dips down a little lower than what would appear to be the handle, and has four vertical slits in the middle allowing air to get through. This rectangular end can easily support and flip light, flat objects. This object thus resembles in all ways what we might refer to as a "spatula." However, this object was found buried in the woods of Eastern Massachusetts where historians are sure Native Americans once lived, and the evidence indicates that this object was most likely used by young boys and men as they constructed habitats for their families.

This object is a spatula. This object is not a spatula. This object is both a spatula and another kind of object. This object is neither a spatula nor another kind of object.

Appendix to Experiment 1

Introduction
Welcome to the study. In this study, you're going to read a series of short stories. After each one, you will be asked some questions. Your job is to answer these questions to the best of your ability, based on what you read. Do not overthink this! This is not a test, and there are no "correct" answers. Upon reading each question, simply answer TRUE or FALSE to each of the statements based on what you read. Let's start out with some stories about some newly-discovered objects or living things. These require a bit of imagination, because they are going to take you into another space and time, and ask you to make some inferences. Just read each story and do your best to answer the questions as best you can.

2. Control items

2.1. Natural kind: Animate

2.1.1. Same Essential, Same Superficial Biologists have recently discovered an animal living in a tropical climate near a water source that, upon inspection, resembles a frog in all of its outward properties.

Preliminary analysis: The animal has eyes protruding from the side of its head, a backbone and a cloaca, webbed hind feet, strong kind legs, smooth moist skin, and a tongue that is sticky and moves incredibly fast. The animal was found among leaves near the water.

Subsequent analysis: The scientists placed food out to examine the animal's diet, and discovered that it will eat insects and spiders, and even larvae and worms. The male lets out a croaking sound to attract the females. The females lay eggs; when hatched, the young start out

looking like legless, fish-like creatures that swim, and then gradually grow legs and emerge onto the land.

This animal is a frog. This animal is not a frog. This animal is both a frog and another kind of animal. This animal is neither a frog nor another kind of animal.

2.1.2. Different Essential, Different Superficial Biologists have recently discovered an animal swimming in the Indian Ocean. Upon inspection, it resembles members of the delphinidae family.

Preliminary analysis: The animal has fins, smooth skin, and plunges to significant depths in the water. It travels in groups of at least 15 and upwards of a thousand. It eats fish and shrimp, and is hunted by sharks.

Subsequent analysis: This animal does not lay eggs. It does not fly or go onto the land. This animal's communication is quite distinct. It can use its body to slap its tail on the surface of the water, or stroke and caress each other. It also has an oil in its head that amplifies sound waves so that it can produce and detect the squeaks and whistles that other animals of its kind emit.

This animal is a bird. This animal is not a bird. This animal is both a bird and another kind of animal. This animal is neither a bird nor another kind of animal.

2.2. Natural kind: Inanimate

2.2.1. Same Essential, Same Superficial Geologists have recently discovered a mineral deposit deep within a test site being excavated in the Antarctic. They brought the mineral back to a government lab in the United States to do an analysis.

Preliminary Analysis: The mineral is hard and crystalline, with protruding tetrahedral spikes. Some samples are clear, while other samples are bright colors, like purple, rose, blue, or yellow. They have determined that on the Mohs hardness scale, a scale used to measure the hardness of minerals, it can be scratched by a masonry drill bit and a steel nail, but not by a glass plate, penny or finger nail, ranking it a 7, alongside quartz.

Subsequent analysis: A more in-depth analysis reveals that its chemical composition is SiO_4 silicon-oxygen, equivalent to quartz.

This mineral is quartz. This mineral is not quartz. This mineral is both quartz and another kind of mineral. This mineral is neither quartz nor another kind of mineral.

2.2.2. Different Essential, Different Superficial A Mining Company has recently adapted mining methods they use here on earth in a way that will allow them to mine for metals on asteroids. On their most recent outer space mining voyage, they recovered some metal samples from an asteroid to bring back to the lab and study further.

Preliminary analysis: The metal has a bright blue color when in a mass, but when finely divided, it appears white. Additionally, the metal appears to be very hard and unmalleable in its pure form, only being able to be scratched by masonry drill bits. This makes it less like gold, which is very soft in its pure form, and more like topaz or diamonds.

Subsequent analysis: A more in-depth analysis of the metal revealed that it has over 150 protons in the atom's nucleus, whereas gold has 79. It does not conduct electricity or heat well, and will corrode easily.

This metal is gold. This metal is not gold. This metal is both gold and another kind of metal. This mineral is neither gold nor another kind of metal.

2.3. Artifact

Function

2.3.1. Same Essential, Same Superficial Archeologists have found an object in a dig in Virginia that has a long wooden handle and a hard metal shape on one end.

Preliminary analysis: The wooden handle appears to have been carved from an oak tree. It is approximately 12" long, and 1" in diameter, and not entirely straight. It fits snugly in a hole in the metal shape. The metal shape is thin and extremely sharp on one end, and on the opposite end, thicker and blunt. The metal is smoothly polished on both sides.

Subsequent analysis: When launched from someone's hands, the object spins through the air, and then lands sharply in a target. If this target were e.g., a watermelon or an animal, it could easily split the target and/or be lodged in it. Based on the objects found in the dig and their placement, this object appears to have been used to cleave wood for burning fires.

This object is an axe. This object is not an axe. This object is both an axe and another kind of object. This object is neither an axe nor another kind of object.

2.3.2. Different Essential, Different Superficial The Egyptian-American Archeological Research Group recently found an object during one of their digs that they have brought back to their lab for further study. During their excavation, they have typically found objects similar in function and appearance to modern day cups, plates, and utensils.

Preliminary analysis: The object is small, cylindrical, and made of clay. At one end of the object's cylindrical shape, it has a concave surface that curves in 3 inches deep, whereas at the other end of the object, the surface is flat and smooth. Thus, the shape and appearance of the object is more similar to a modern day cup than it is, say, to a plate or a utensil.

Subsequent analysis: Upon learning that the excavation site where this object was recovered from was an ancient brewery, the researchers also discovered that the ancient Egyptians would use these small objects as a container for holding and drinking beer.

This object is a plate. This object is not a plate. This object is both a plate and another kind of object. This object is neither a plate nor another kind of object.

Intention

2.3.3 Same Essential, Same Superficial A product review company was hired to conduct an analysis of some products made by the Northeastern Manufacturing Corporation. To begin the process, the review company was first asked to examine one particular object.

Preliminary analysis: The shape and appearance of the object resembled an ordinary bookshelf. It had four pieces of wood that were connected together in a rectangular box shape with a flat, thin piece of wood covering one side of the box. On the inside of this box shape, there were two horizontal pieces of wood that were connected to the long edges of the box.

Subsequent analysis: Upon interviewing the employee who built this object, the review company learned that she intended it to be a bookshelf.

This object is a bookshelf. This object is not a bookshelf. This object is both a bookshelf and another kind of object. This object is neither a bookshelf nor another kind of object.

2.3.4 Different Essential, Different Superficial The American Soda Corporation recently hired an employee to manufacture bottle caps for their soda products. Additionally, to evaluate the manufacturing employee's performance on this task, they also hired an external review company to conduct an analysis of the products the employee made.

Preliminary analysis: The object was small, cylindrical, and made of cork (bark of the cork

oak). The shape and appearance of the object seemed to resemble a solid wine cork more than a bottle cap.

Subsequent analysis: Upon interviewing the manufacturing employee, the review company learned that he intended it to be a wine cork and not a bottle cap.

This object is a bottle cap. This object is not a bottle cap. This object is both a bottle cap and another kind of object. This object is neither a bottle cap nor another kind of object.

Appendix to Experiment 2

Introduction

Now let's move on to some new stories. These stories are about an inventor who has built a new machine. Here's how it works: the machine takes as input everyday objects, performs some alterations or steps, and produces as output everyday objects.

Sometimes the alterations are extremely minimal (perhaps changing nothing at all). At other times, they are quite drastic (changing quite a lot!). After the machine performs this operation, the inventor presents the output-objects to a focus group to assess the appearance and functionality of the output-object and to see how they compare to the input-objects in this regard.

2.1. Same essential, different superficial

2.1.1. Teakettle/birdhouse The inventor fed a tea kettle into the machine as input. First, the machine sawed off the handle, made the spout shorter, spray painted the surface pink, and removed the knob on the lid. After this, the machine attached a flat piece of metal that protruded out from the bottom and a pointy covering on the top. Finally, the machine added a short stick that stuck out from the side of the object and cut a small hole in the middle of the side of the object just above the stick. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it now had the outward appearance of a birdhouse but determined it could not function as such. However, while the object did not have the original appearance of a tea kettle, it still had the capability to heat and pour hot water.

After the operation, this object was a tea kettle. After the operation, this object was a birdhouse. After the operation, this object was both a tea kettle and a birdhouse. After the operation, this object was neither a tea kettle nor a birdhouse.

2.1.2. Garbage can/chair The inventor fed a garbage can into the machine as input. First, the machine spray painted the surface orange and attached four cylindrical pieces of wood to the bottom of the object. After this, the machine attached a soft removable cushion to the top of the lid. Finally, the machine attached a flat piece of metal to the side of the object perpendicular to the floor, which protruded three feet past the top of the lid. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it now had the outward appearance of a chair but determined that it could not function as such. However, while the object did not have the original appearance of a garbage can, it could still be used as a container for storing waste.

After the operation, this object was a garbage can. After the operation, this object was a chair. After the operation, this object was both a garbage can and a chair. After the operation,

this object was neither a garbage can nor a chair.

2.1.3. PVC pipe/flute The inventor fed a PVC pipe into the machine as input. First, the machine spray painted the surface a metallic silver color. After this, the machine drilled sixteen small holes in a straight line on the surface of the object. Finally, the machine mounted twelve small, round, padded metal keys to cover and uncover the drilled holes and attached a mouthpiece at one end of the object. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it now had the outward appearance of a flute, but determined it could not function as such. However, while the object did not have the original appearance of a PVC pipe, it could still be used as a container for various piping applications such as transporting drinking water over drainage systems or servicing sprinkler systems.

After the operation, this object was a PVC pipe. After the operation, this object was a flute. After the operation, this object was both a PVC pipe and a flute. After the operation, this object was neither a PVC pipe nor a flute.

2.1.4. Cards/wipes The machine took a deck of 52 playing cards as input. First, it bleached the top side of them. Then it applied a special liquid to soften them so that they were much more pliable. It then applied considerable pressure to the stack of cards, pressing them down so that they became much wider yet still rectangular. Finally, the machine injected each card with a substance so that it was slippery to the touch and droopy when lifted. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it now had the outward appearance of a stack of wet wipes, but determined it could not function as such. However, while the object did not have the original appearance of a deck of cards, it could still be used to play games like Old Maid and Poker.

After the operation, this object was a deck of cards. After the operation, this object was a stack of wet wipes. After the operation, this object was both a deck of cards can and a stack of wet wipes. After the operation, this object was neither a deck of cards nor a stack of wet wipes.

2.2. Same essential, same superficial 2.2.1. Teakettle/birdhouse The inventor fed a tea kettle into the machine as input. First, the machine sawed off the handle, then re-attached it, cut the spout shorter, then attached an extension to reach the original length, sanded the surface, then applied a primer and the original color, removed the knob on the lid and polished it and then re-attached it. After this, the machine attached a flat piece of metal that protruded out from the bottom and a pointy covering on the top. It spun the object around, and then removed this metal base and the pointy covering. Finally, the machine added a short stick that stuck out from the side of the object and cut a small hole in the middle of the side of the object just above the stick. It then popped off the stick and plugged up the hole. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a tea kettle and also could be used to heat and pour hot water.

After the operation, this object was a tea kettle. After the operation, this object was another object. After the operation, this object was both a tea kettle and another object. After the operation, this object was neither a tea kettle nor another object.

2.2.2. Garbage can/chair The inventor fed a garbage can into the machine as input. First,

the machine buffed the surface of the object and then polished it. Then it attached four cylindrical pieces of wood to the bottom of the object, turned the object 360° and then removed the cylinders. After this, the machine attached a soft removable cushion to the top of the lid, opened and closed it three times, and then removed the cushion. Finally, the machine attached a flat piece of metal to the side of the object perpendicular to the floor, which protruded three feet past the top of the lid, stuck and unstuck this piece five times, and then removed it. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a garbage can and could be used as a container for storing waste.

After the operation, this object was a garbage can. After the operation, this object was another object. After the operation, this object was both a garbage can and another object. After the operation, this object was neither a garbage can nor another object.

2.2.3. PVC pipe/flute The inventor fed a PVC pipe into the machine as input. First, the machine sanded down the pipe, then buffed and polished the surface. After this, the machine drilled sixteen small holes in a straight line on the surface of the object, and then mounted twelve small, round, padded metal keys to cover and uncover the drilled holes. It pressed up and down on these keys five times each before removing them and then plugging up the holes securely and air tight. Finally, it attached a mouthpiece at one end of the object, twisted this attachment clockwise and then counter clockwise, and then removed it. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a PVC pipe and could be used as a container for various piping applications such as transporting drinking water over drainage systems or servicing sprinkler systems.

After the operation, this object was a PVC pipe. After the operation, this object was another object. After the operation, this object was both a PVC pipe and another object. After the operation, this object was neither a PVC pipe nor another object.

2.2.4. Cards/wipes The machine took a deck of 52 playing cards as input. First, it shuffled the cards 20 times. Then it applied a washable stamp to each of their sides, and then removed this with a quick-drying chemical. It then blew short bursts of air at the object from all sides. Finally, it removed a card out of the middle and placed it on the top. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a deck of cards and could also be used to play games like Old Maid and Poker.

After the operation, this object was a deck of cards. After the operation, this object was another object. After the operation, this object was both a deck of cards and another object. After the operation, this object was neither a deck of cards nor another object.

2.3. Different essential, same superficial

2.3.1. Teakettle/birdhouse The inventor fed a tea kettle into the machine as input. First, the machine sawed off the handle, then re-attached it, cut the spout shorter, then attached an extension to reach the original length, sanded the surface, then applied a primer and the original color, removed the knob on the lid and polished it and then re-attached it. After this, the machine attached a flat piece of metal that protruded out from the bottom and a pointy covering on the top. It spun the object around, and then removed this metal base and the

pointy covering. Finally, the machine added a short stick that stuck out from the side of the object and cut a small hole in the middle of the side of the object just above the stick. It then popped off the stick and plugged up the hole. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a tea kettle. At the same time, it no longer had the capability to heat and pour hot water.

After the operation, this object was a tea kettle. After the operation, this object was another object. After the operation, this object was both a tea kettle and another object. After the operation, this object was neither a tea kettle nor another object.

2.3.2. Garbage can/chair The inventor fed a garbage can into the machine as input. First, the machine buffed the surface of the object and then polished it. Then it attached four cylindrical pieces of wood to the bottom of the object, turned the object 360° and then removed the cylinders. After this, the machine attached a soft removable cushion to the top of the lid, opened and closed it three times, and then removed the cushion. Finally, the machine attached a flat piece of metal to the side of the object perpendicular to the floor, which protruded three feet past the top of the lid, stuck and unstuck this piece five times, and then removed it. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a garbage can, but determined that it could no longer be used as a container for storing waste.

After the operation, this object was a garbage can. After the operation, this object was another object. After the operation, this object was both a garbage can and another object. After the operation, this object was neither a garbage can nor another object.

2.3.3. PVC pipe/flute The inventor fed a PVC pipe into the machine as input. First, the machine sanded down the pipe, then buffed and polished the surface. After this, the machine drilled sixteen small holes in a straight line on the surface of the object, and then mounted twelve small, round, padded metal keys to cover and uncover the drilled holes. It pressed up and down on these keys five times each before removing them and then plugging up the holes securely and air tight. Finally, it attached a mouthpiece at one end of the object, twisted this attachment clockwise and then counter clockwise, and then removed it. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a PVC pipe, but determined it could not function as such, since it could no longer be used as a container for various piping applications such as transporting drinking water over drainage systems or servicing sprinkler systems.

After the operation, this object was a PVC pipe. After the operation, this object was another object. After the operation, this object was both a PVC pipe and another object. After the operation, this object was neither a PVC pipe nor another object.

2.3.4. Cards/wipes The machine took a deck of 52 playing cards as input. First, it shuffled the cards 20 times. Then it applied a washable stamp to each of their sides, and then removed this with a quick-drying chemical. It then blew short bursts of air at the object from all sides. Finally, it removed a card out of the middle and placed it on the top. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a deck of cards, but determined it could not function as such, since it could no longer be used to play games like Old Maid and Poker.

After the operation, this object was a deck of cards. After the operation, this object was another object. After the operation, this object was both a deck of cards and another object. After the operation, this object was neither a deck of cards nor another object.

2.4. Different essential, different superficial

2.4.1. Teakettle/birdhouse An inventor has created a machine that takes as input everyday objects, performs some alterations, and produces as output everyday objects. Sometimes the alterations are extremely minimal; other times, they are quite drastic. The inventor fed a tea kettle into the machine as input. First, the machine sawed off the handle, made the spout shorter, spray painted the surface pink, and removed the knob on the lid. After this, the machine attached a flat piece of metal that protruded out from the bottom and a pointy covering on the top. Finally, the machine added a short stick that stuck out from the side of the object and cut a small hole in the middle of the side of the object just above the stick. When the machine finished its operation on the object, the inventor presented it to a focus group of people, who indicated that it now had the outward appearance of a birdhouse but determined it could not function as such. Moreover, while the object did not have the original appearance of a tea kettle, it also no longer had the capability to heat and pour hot water.

After the operation, this object was a tea kettle. After the operation, this object was a birdhouse. After the operation, this object was both a tea kettle and a birdhouse. After the operation, this object was neither a tea kettle nor a birdhouse.

2.4.2. Garbage can/chair The inventor fed a garbage can into the machine as input. First, the machine spray painted the surface orange and attached four cylindrical pieces of wood to the bottom of the object. After this, the machine attached a soft removable cushion to the top of the lid. Finally, the machine attached a flat piece of metal to the side of the object perpendicular to the floor, which protruded three feet past the top of the lid. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it now had the outward appearance of a chair but determined that it could not function as such. Moreover, while the object did not have the original appearance of a garbage can, it also could no longer be used as a container for storing waste.

After the operation, this object was a garbage can. After the operation, this object was a chair. After the operation, this object was both a garbage can and a chair. After the operation, this object was neither a garbage can nor a chair.

2.4.3. PVC pipe/flute The inventor fed a PVC pipe into the machine as input. First, the machine spray painted the surface a metallic silver color. After this, the machine drilled sixteen small holes in a straight line on the surface of the object. Finally, the machine mounted twelve small, round, padded metal keys to cover and uncover the drilled holes and attached a mouthpiece at one end of the object. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it now had the outward appearance of a flute, but determined it could not function as such. Moreover, while the object did not have the original appearance of a PVC pipe, it could also no longer be used as a container for various piping applications such as transporting drinking water over drainage systems or servicing sprinkler systems.

After the operation, this object was a PVC pipe. After the operation, this object was a flute. After the operation, this object was both a PVC pipe and a flute. After the operation, this object

was neither a PVC pipe nor a flute.

2.4.4. Cards/wipes The machine took a deck of 52 playing cards as input. First, it bleached the top side of them. Then it applied a special liquid to soften them so that they were much more pliable. It then applied considerable pressure to the stack of cards, pressing them down so that they became much wider yet still rectangular. Finally, the machine injected each card with a substance so that it was slippery to the touch and droopy when lifted. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it now had the outward appearance of a stack of wet wipes, but determined it could not function as such. Moreover, while the object did not have the original appearance of a deck of cards, it could also no longer be used to play games like Old Maid and Poker.

After the operation, this object was a deck of cards. After the operation, this object was a stack of wet wipes. After the operation, this object was both a deck of cards can and a stack of wet wipes. After the operation, this object was neither a deck of cards nor a stack of wet wipes.

Appendix to Experiment 3

Introduction

Welcome to the study! In this study, you're going to read a series of short stories. After each one, you are going to see two statementA: one by Speaker A and another by Speaker B. These are two different speakers.

We are then going to ask you to read a set of four statements and select the one with which you most agree.

The four statements following each story will be the same:

A is RIGHT and B is WRONG.

A is WRONG and B is RIGHT.

There is a sense in which both A and B are RIGHT.

Neither A nor B is right.

Simply go with your gut reaction and pick the statement with which you most agree. There is no right or wrong answer.

Remember, do not overthink this! This is not a test, and there are no "correct" answers! However, please note that you will NOT get credit for patterning randomly or taking less than 15 minutes to complete the study. So please read the stores carefully and make a principled decision.

Let's begin!

2. Control Items

2.1. PPT: FD acceptable

2.1.1. Story 1: New salsa A local Tex Mex restaurant has introduced a new salsa to their menu. It is supposed to be an alternative to their extremely mild pico de gallo and their extremely spicy chipotle sauce. Customers have been flocking to the restaurant this week to try the new salsa on their burritos, chips, and quesadillas, and have been sharing their reactions on social media. A and B could not miss out on this experience. They go to the restaurant to try the salsa, and report the following.

A: This salsa is spicy. B: This salsa is not spicy.

2.1.2. Story 2: Baseball game

A and B went to watch a local baseball team play their out-of-town rivals. Both teams are ranked near the middle of their conference, but both have some outstanding players who made some clutch plays throughout the game. The score remained 0-0 until the 7th inning, when each team finally scored a run and tied the game. The home team hit a game-winning home run in extra innings and won with a score of 2-1. As they are leaving the baseball stadium to go to a pub, A and B report the following.

A: That baseball game was fun. B: That baseball game was not fun.

2.2. Relative GA: FD acceptable

2.2.1. Story 1: Guy at picnic There was a company picnic today at a local park, which allowed for everyone in attendance to get some fresh air, enjoy the sunlight, and participate in team building activities and icebreakers. E and S are coworkers in attendance. As they are talking to each other, they spot a guy some distance away, who neither of them knows, and is serving himself some potato salad. A is 6'2" and B is 5'2". The guy is 5'10". As A and B are talking and are watching him serve himself, they report the following.

A: That guy is not tall. B: That guy is tall.

2.2.2. Story 2: Urban latte

A and B are in the city for the day to attend a conference for work and have some meetings with clients. They walk into an urban coffee shop so they can each order a latte. A latte at this coffee shop costs $5.25. A is an administrative assistant who gets paid by the hour and B is an executive with a steady salary. As A and B are looking at the coffee prices on the wall, they report the following.

A: A latte at this coffee shop is expensive. B: A latte at this coffee shop is not expensive.

2.3. FD not acceptable

2.3.1. Story 1: Portland A and B are planning on going on a road trip this summer, and are sharing ideas about where they might travel. They have both been to Maine, Virginia, and Texas. They are not as familiar with the Western part of the US. Someone they both know has recommended that they visit Portland. While A and B are considering this as a destination, they report the following.

A: Portland is a state. B: Portland is a city.

2.3.2. Story 2: Apples A and B are in the produce section of the grocery store. They are talking about eating healthy and what kind of food to buy. They agree that it is best to eat lots of leafy greens, fruits, and vegetables. As they talk about what kinds of foods to purchase and how to prepare them, A and B report the following.

A: An apple is a fruit. B: An apple is a vegetable.

2.3.3. Story 3: Quantity A and B like to do fun challenges, especially if it involves eating. They have recently decided to try a new ice cream flavor for every week of the year. They are trying to figure out how many flavors they will be exploring over the course of a year. As they look ahead to all the flavors they will be trying, A and B report the following.

A: There are 45 weeks in a year. B: There are 52 weeks in a year.

2.3.4. Story 4: Date A and B enjoy Taco Tuesday every week without fail. Two days ago, bright and early on a Sunday morning when they did their cardio kick-boxing exercise session

in the park together, they reminded each other of what they could soon look forward to that week. Just today when they were talking on the phone, A and B reported the following.

A: Today is Tuesday B: Today is not Tuesday.

3. Control: Transformation Introduction Now let's move on to some new stories. These stories will be about an inventor who has built a new machine. Here's how it works: the machine takes as input everyday objects, performs some alterations or steps, and produces as output everyday objects.

Sometimes the alterations are extremely minimal (perhaps changing nothing at all). At other times, they are quite drastic (changing quite a lot!). After the machine performs this operation, the inventor presents the output-objects to a focus group to assess the appearance and functionality of the output-object and to see how they compare to the input-objects in this regard.

Your task is still the same! Upon reading each question, simply choose the statement with which you most agree.

2.1. Same essential, different superficial: Teakettle/birdhouse The inventor fed a tea kettle into the machine as input. First, the machine sawed off the handle, made the spout shorter, spray painted the surface pink, and removed the knob on the lid. After this, the machine attached a flat piece of metal that protruded out from the bottom and a pointy covering on the top. Finally, the machine added a short stick that stuck out from the side of the object and cut a small hole in the middle of the side of the object just above the stick. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it now had the outward appearance of a birdhouse but determined it could not function as such. However, while the object did not have the original appearance of a tea kettle, it still had the capability to heat and pour hot water.

A: After the operation, this object was another kind of object. B: After the operation, this object was a tea kettle.

2.2. Same essential, same superficial: PVC pipe/flute The inventor fed a PVC pipe into the machine as input. First, the machine sanded down the pipe, then buffed and polished the surface. After this, the machine drilled sixteen small holes in a straight line on the surface of the object, and then mounted twelve small, round, padded metal keys to cover and uncover the drilled holes. It pressed up and down on these keys five times each before removing them and then plugging up the holes securely and air tight. Finally, it attached a mouthpiece at one end of the object, twisted this attachment clockwise and then counter clockwise, and then removed it. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a PVC pipe and could be used as a container for various piping applications such as transporting drinking water over drainage systems or servicing sprinkler systems.

A: After the operation, this object was a PVC pipe. B: After the operation, this object was another kind of object.

2.3. Different essential, same superficial: Garbage can/chair The inventor fed a garbage can into the machine as input. First, the machine buffed the surface of the object and then polished it. Then it attached four cylindrical pieces of wood to the bottom of the object, turned the object 360° and then removed the cylinders. After this, the machine attached a soft removable cushion

to the top of the lid, opened and closed it three times, and then removed the cushion. Finally, the machine attached a flat piece of metal to the side of the object perpendicular to the floor, which protruded three feet past the top of the lid, stuck and unstuck this piece five times, and then removed it. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it had the outward appearance of a garbage can, but determined that it could no longer be used as a container for storing waste.

A: After the operation, this object was a garbage can. B: After the operation, this object was another kind of object.

2.4. Different essential, different superficial: Cards/wipes The machine took a deck of 52 playing cards as input. First, it bleached the top side of them. Then it applied a special liquid to soften them so that they were much more pliable. It then applied considerable pressure to the stack of cards, pressing them down so that they became much wider yet still rectangular. Finally, the machine injected each card with a substance so that it was slippery to the touch and droopy when lifted. When the machine finished its operation on the object, the inventor presented it to a focus group of people who indicated that it now had the outward appearance of a stack of wet wipes, but determined it could not function as such. Moreover, while the object did not have the original appearance of a deck of cards, it could also no longer be used to play games like Old Maid and Poker.

A: After the operation, this object was a deck of cards. B: After the operation, this object was another kind of object.