

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

“Jack is a True Scientist”: On the Content of Dual Character Concepts

#### **Permalink**

<https://escholarship.org/uc/item/5c20g816>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

#### **Authors**

Del Pinal, Guillermo

Reuter, Kevin

#### **Publication Date**

2015

Peer reviewed

# “Jack is a True Scientist”: On the Content of Dual Character Concepts

Guillermo Del Pinal (ged2102@columbia.edu)

Department of Philosophy, Columbia University, USA  
Department of Philosophy, Ruhr University Bochum, Germany

Kevin Reuter (kevin.reuter@rub.de)

Department of Philosophy, Ruhr University Bochum, Germany

## Abstract

The concepts expressed by social role terms such as *artist* and *scientist* are unique. In a series of experiments, Knobe, Prasada, and Newman (2013) show that these terms allow two independent criteria for categorisation, one of which is inherently normative. This paper presents and tests a novel account of the content of these ‘dual character concepts’. We argue that the normative dimension of dual character concepts represents commitments to fulfill certain idealized functions. We then present evidence that the normative dimension is a central dimension in the conceptual structure of dual character concepts. Finally, we show that our account is both descriptively and explanatorily adequate.

**Keywords:** dual character concepts; social cognition; social role concepts.

## 1 Introduction

In a recent paper, Knobe, Prasada, and Newman (2013) present a series of original experiments designed to show that the concepts expressed by terms such as *artist* and *scientist* have two independent criteria for categorization, one of which is inherently normative. They call this unique class of concepts ‘dual character concepts’ (DCCs). The present work presents and defends an account of the content of DCCs. To illustrate the basic properties of DCCs consider the following scenarios. John has two biological children. He materially provides them with all their needs, including food, proper schooling, and some nice toys. However, John does not think doing so is his duty; in fact, he is only a good father because he thinks that his behavior will advance his career as a politician. Is John really a father? In response, consider whether you agree with (1-a) and (1-b):

- (1) a. There is a sense in which John is clearly a father.  
b. However, if you think about what it really means to be a father, you would have to say that John is not a true father after all.

Now, imagine that John’s career foundered and that he abandoned his two kids. Fortunately, John’s brother, Mark, who has no biological children, decided to care for the kids. Mark is not as materially successful as John. Often he can’t provide the kids with some of their needs. Still, Mark really loves the kids and works very hard to make sure they have everything they need. Is Mark really a father? In response, consider whether you agree with (2-a) and (2-b):

- (2) a. There is a sense in which Mark is clearly not a father.  
b. However, if you think about what it really means to be a father, you would have to say that Mark is a true father after all

The experiments presented by Knobe et al. (2013) suggest that most of us would accept these statements, and that this basic pattern generalizes to other social role terms such as *artist* and *scientist*. Now, not all social roles are equally acceptable in this basic dual pattern scenario. In particular, terms such as *doorman* and *welder* obtain lower acceptability ratings in cases parallel to (1-b) and (2-b).

Knobe et al. argue that to explain these acceptability patterns, we need to assume that DCCs have a unique internal structure. They agree with the traditional view that we often use concrete or salient features to categorize objects, even when we know that some entities might (i) have these features and not fall under the category, or (ii) lack these features and still fall under the category (Lakoff, 1973; Gelman, 2003). For example, we know that some substance might look exactly like gold and not be gold, and that some animal that lacks the typical tiger coat might still be a tiger. But what distinguishes, in the first instance, DCCs from other kinds of lexical concepts is that DCCs include a normative dimension on which modifiers such as *true* operate.

What is the content of the normative dimension? The experiments of Knobe et al. partially address that question, but a full answer requires additional refinements and experimental tests. In what follows, we present an account of DCCs that explicitly addresses this question. Specifically, we will test and defend the following hypotheses: (i) the normative dimension represents the commitment to fulfil the idealized functions associated with DCCs, and (ii) this dimension is central to DCCs. On this view, someone counts as a true father not because he actually fulfills the idealized function of a father, but because he is committed to doing so, even if he happens to fail. In addition, this commitment is a central dimension of the category of fathers, i.e., it is a dimension which explains the inclusion of other concrete features associated with fathers.

In §2 we present a preliminary study used to generate stimuli for our experiments. In §3 we present and in §4 we test our account. In §5 we argue that our theory of DCCs is explanatorily adequate.

## 2 Preliminary Study

Our study of DCCs focuses on social role terms such as *firefighter*, *bartender*, and *scientist*. Broadly speaking, these terms pick out professional social roles. Knobe, Prasada, and Newman (2013) argue that one way to determine whether a term expresses a DCC is to consider its acceptability under *true*-modification. By ‘*true*-modification’ we mean sentences of the form of (3), i.e., in which the predicate is modified by *true*:

(3) Jack is a true scientist.

The basic idea is that, in this kind of linguistic construction, the *true*-modifier operates on the normative dimension. *Scientist* is a paradigmatic DCC, hence (3) seems perfectly acceptable. For terms that arguably have either no normative dimension or no default value on it, *true*-modifications will seem less acceptable, as in (4):

(4) Jack is a true cashier.

As this example illustrates, not all social role terms are equally acceptable under *true*-modification. Whether this means that only some social role terms are really DCCs—as Knobe et al. (2013) and Leslie (2014) seem to think—is an issue that we shall discuss later. For now, let us simply call social role terms which receive low acceptability ratings in *true*-modifications ‘weak DCCs’ and those which receive high-acceptability ratings ‘strong DCCs’. Although we should distinguish between weak and strong DCCs, it is also important to note that DCCs are highly context sensitive (Leslie 2014). More precisely: adding a context can increase the acceptability ratings of *true*-modifications involving social role terms that, in other conditions, receive low ratings. This suggests that one way of investigating the content of DCCs is to examine the interaction between contextual parameters and shifts in the acceptability of weak and strong DCCs in *true*-modifications. We follow this strategy in Study 1. Given this aim, we need a set of stimuli that allows us to trace and compare the behavior of weak and strong DCCs. The aim of the following preliminary study is to generate such a set. All subjects who participated in our experiments were recruited through Amazon’s Mechanical Turk system for human intelligence tasks.

### Method

276 subjects participated in this preliminary study. Each participant received a single question of the form *Jack is a true x* without any context. Following Knobe et al. (2013)’s measure, participants were then asked to rate how natural/weird this sentence sounded on a 7-point Likert scale from ‘1’ meaning ‘sounds weird’ to ‘7’ meaning ‘sounds natural’. We used all the professional-role concepts that in Knobe et al. (2013)’s original study were classified as either DCCs or controls.

### Results and Discussion

The average ratings for each individual concept are shown in Figure 1. The mean rating for those concepts that Knobe et

al. classified as DCCs is 5.52 (SE = .16) and for controls is 4.71 (SE = .15). An independent t-test was performed,  $t(256) = 3.743$ ;  $p < .001$ , indicating that the difference between DCCs and social role controls is highly significant.




Figure 1: Acceptability-ratings for *true*-modifications of 18 social role concepts.

This preliminary study replicates a key result of Knobe et al. (2013). When judging the acceptability of *true*-modifications, the mean rating for DCCs such as *artist* is significantly higher than the mean rating for the social role controls such as *doorman*. At the same time, if we list the individual ratings, as in Figure 1, we can see that there is a smooth transition between the acceptability ratings of the highest and lowest rated DCCs. This suggests that we should talk instead of strong and weak DCCs. To examine and compare the effect of context on the ratings of DCCs in our subsequent experiment, we used a median split dividing the concepts into strong and weak DCCs.

## 3 The normative dimension of DCCs

What is the content of the normative dimension of DCCs? When we say that the concept expressed by *scientist* includes a ‘normative component’ and not only a set of concrete features, what exactly are we claiming is represented? Knobe et al. (2013) raise some key points. First, this dimension has to do with abstract values. What they mean by this is revealed in their examples: for *scientist* the abstract value is something like the quest for impartial truths and explanations. Second, the normative dimension of a class must be distinguished from its typical function. The normative dimension seems to directly interact with modifiers such as *true* but not directly with modifiers such as *good*, taken in the sense of ‘useful or efficient’ (see Knobe et al. (2013)’s experiment 1). In the intended reading, a *good calculator* is an artifact that is good for its function; similarly for social roles such as *efficient* or *skillful scientist*. However, we can imagine scenarios in which someone can be a true scientist without being, yet, a skillful scientist. We can also imagine scenarios in which someone can be a skillful scientist without, ultimately, being a true scientist (see study 1 in §4.1). Since, in the intended reading, the acceptability of *good/skillful* vs. *true scientist* can dissociate, it seems clear that the normative dimension does not directly represent the typical function of the class.

Knobe et al. (2013) do not commit to a more specific

claim about the content of the normative dimension; but building on their results, Leslie (2014) makes a concrete proposal. Leslie agrees that the normative dimension of DCCs does not represent their typical function, but she argues that it represents something that is closely related, namely, their ‘ideal function’. On this view, we can say that a scientist that completely resists empirical demands to modify or abandon his theory is not a true scientist because he is not satisfying the ideal function of a scientist, which is something like building theories that are responsive to empirical evidence, among other things. Leslie’s proposal can be implemented in different ways. One way is to assume that the idealized function is represented directly. On this view, the normative dimension of *scientist* would represent something like the search for empirically constrained truths and theories. However, this is not the best way to implement Leslie’s proposal. For it is in tension with the basic observation, mentioned above, according to which someone who tries hard but often fails to embody the ideal function of a father or scientist can still count as a true father or scientist. Study 1 below confirms this basic intuition.

A better way to implement Leslie’s proposal, in our view, is to hold that what really counts to satisfy the normative dimension is not so much whether someone actually fulfills the idealized function to some non-trivial degree, but rather whether someone is committed to fulfill the idealized function. Consider an example of non-professional or not yet skillful scientist, Jim. On this view, we can still say that Jim is a true scientist as long as he is clearly committed to the ideals of producing theories and views that are responsive to evidence. So according to this account, there are many contexts in which someone who fails as a scientist nevertheless counts as a *true* scientist, as long as s/he has the required commitments.

Our hypothesis then is that we can dissociate the property of actually fulfilling the idealized function from that being committed to fulfill the idealized function, and that what ultimately matters for satisfying the normative dimension is the commitment. Study 1 below tests and supports this view of the relation between commitments, idealized functions and the content of the normative dimension of DCCs. To be clear, our claim is not that evidence that an entity fulfills the idealized function does not matter for satisfying the normative dimension of DCCs. After all, that someone is a really an efficient or skillful scientist is often a reliable sign that s/he is committed to being a scientist. Our claim is that, when pitted against each other, being committed to fulfill the idealized function is more important for satisfying the normative dimension of DCCs than simply being able to fulfill that function.

## 4 Experimental Studies

### 4.1 Study 1

We have argued that the normative dimension does not directly represent typical or ideal functions, but rather represents these as the objects of commitments. What is the evidence for this? The vignettes used by Knobe et al. (2013)

in their Experiment 2 (the most relevant for our purposes) do not assert that the relevant individuals are good or skillful at their role; hence they cannot be used to support the alleged dissociation between the functional and the normative dimensions. To directly test the hypothesis that what matters for the normative dimension is not whether one is skilled or efficient, but rather whether one is committed, we presented subjects with ‘high-function’ and ‘high-devotion’ contexts and compared their effects on the acceptability of *true*-modifications. The high-function context highlights Jack’s high skill as an x. We were particularly interested to see whether an explicit high-function context would increase the acceptability ratings of *true*-modifications of weak DCCs. The high-devotion context highlights Jack’s high devotion to being an x. If our view is correct, the high-devotion context should have a stronger and more positive impact on the acceptability of *true*-modifications than the high-function context.

### Method

We used the 18 social role concepts tested in the preliminary study. Based on those results, we did a median split, dividing them into weak DCCs (*caseworker, doorman, cashier, waiter, tailor, teacher, minister, optician, bartender*) and strong DCCs (*baker, welder, firefighter, scientist, artist, mechanic, comedian, musician, soldier*). To examine how high-function (HF) and high-devotion (HD) contexts interact with the normative dimension, we presented 271 subjects with HF contexts and a different group of 294 subjects with HD contexts. The respective contexts read as follows:

(HF) Jack is an artist/doorman. He is really skilled and highly efficient at his job.

(HD) Jack is an artist/doorman. He really likes his job and is completely devoted to it.

After being presented with a HF or HD context, participants were then asked to rate whether the sentence *Jack is a true artist/doorman* sounded weird (1), natural (7), or anything in between. We computed the mean values of people’s responses given HF and HD contexts, and compared them with the mean values in the no-context condition (see preliminary study).




Figure 2: Acceptability-ratings for strong and weak DCCs in the no-context condition (left), in the high function context (middle), and in the high devotion context (right).

## Results

The average ratings for strong and weak DCCs in all three conditions are shown in Figure 2. We applied a 3 x 2 ANOVA with participants' ratings as the dependent measure and the independent factors *Context* with three levels (no-context, HF, HD) and *Concept* (strong DCCs, weak DCCs). The analysis reveals highly significant main effects of both *Context*,  $F(2; 835) = 10.571$ ;  $p < 0.001$ , and *Concept*,  $F(835) = 46.227$ ;  $p < 0.001$ . For all 18 social role concepts, the HD context raises the mean value from 5.04 (no-context) to 5.37 (SE = 0.10), whereas the HF context decreases the value to 4.68 (SE = 0.11). Moreover, a highly significant interaction occurs between *Context* and *Concept*,  $F(2; 835) = 4.807$ ;  $p = 0.008$ . For strong DCCs, the HF context decreases the participants' ratings from 5.74 to 5.04, while the HD context has no effect on people's ratings. In contrast, for weak DCCs, the HF context has no effect, whereas the HD context strongly increases their average value from 4.38 (no-context) to 5.15.

## Discussion

The comparative influence of high-function and high-devotion contexts on strong and weak DCCs strongly supports our hypothesis that the normative dimension of DCCs represents the commitments to fulfill the relevant function. Paradigmatic representatives of strong DCCs are terms such as *scientist* and *artist*. Recall that, as established in our preliminary study, strong DCCs are those DCCs that already obtain high acceptability ratings under *true*-modification in the no-context condition. In other words, knowing that some random Jack is a scientist suggests to participants, even without further information, that Jack is likely also committed to the role of a scientist. In contrast, weak DCCs such as *doorman* and *cashier* are those DCCs that obtained lower ratings in the no context condition. In other words, merely knowing that Jack is a cashier does not suggest to participants, to the same degree as with strong DCCs, that Jack satisfies the respective commitments. Unsurprisingly, then, a high-devotion context is less informative, and hence has almost no effect, on the acceptability ratings of strong DCCs. Crucially, however, a high-devotion context basically turns weak DCCs into strong DCCs.

Our results also show that high function contexts decrease the acceptability of *true*-modifications. This further supports the view that the normative dimension of DCCs does not directly represent the typical functions. In particular, even though weak DCCs have space to increase their ratings under *true*-modification, high function contexts did not increase their acceptability. To illustrate, knowing that Jack is a skilled waiter does not increase the acceptability of the claim that Jack is a true waiter. What about the negative effect of high function contexts on the acceptability ratings of strong DCCs? This is probably due to pragmatic effects. Asserting that Jack is a skilled artist can plausibly be taken as evidence that he is at most only a skilled artist, hence that he lacks something which is relevant to his being a true artist. Overall, Study 1 shows that a high devotion context increases the acceptability of

*true*-modifications for weak DCCs. In addition, high function contexts have a negative influence on strong DCCs, and a lack of effect on the acceptability of weak DCCs. Together, these results support the hypothesis that the normative dimension of DCCs represents commitments to fulfill the function of social roles.

## 4.2 Study 2

This study was designed to directly compare the idealized function and the commitment dimensions represented by DCCs. Instead of investigating social role concepts under *true*-modifications, we wanted to compare, in terms of their corresponding centrality, the function and the commitment to the function in the structure of DCCs. To do this, we had to select a measure of conceptual centrality. Sloman et al. (1998) show that there is a direct correlation between the centrality and the immutability of features. The immutability of a feature or dimension reflects how much the internal structure of a concept depends on that feature or dimension. The basic idea is simple. The centrality of a feature in a conceptual representation *C* reflects the degree to which other features in *C* depend on it, i.e., the degree to which it lends conceptual coherence. This means that it should be easier to mentally transform a non-central compared to a central feature of *C*. For example, it should be easier to think of a gun that is neither black nor made of steel, than to think of a gun that is not made to shoot. This is what we mean by the 'mutability' of a feature. As our measure of mutability, hence of conceptual centrality, we chose a paradigm introduced by Sloman et al. (1998), called 'surprise'. In this paradigm, experimenters ask participants to rate how surprised they would be to encounter an instance of the category that does not have some feature. The basic idea is that 'instances missing mutable features should be less surprising than instances missing immutable features on the assumption that surprise is related to the difficulty of adapting an object representation to a concept. Adaptation should be easy if the object is missing a mutable feature but hard if it is missing an immutable feature' (Sloman et al., 1998)

## Method

117 subjects took part in experiment 3. Each participant was asked to evaluate one question that was randomly selected from the following two categories: (1) Commitment to fulfill Idealized Function, e.g. How surprised would you be to encounter an artist who does not care about creating inspiring works of art?; (2) Idealised Function, e.g. How surprised would you be to encounter an artist who does not create inspiring works of art? Subjects were asked to answer the question on a 7 point Likert scale from 1 = 'not at all surprised' to 7 = 'very surprised'. We investigated the social role concepts *artist*, *musician*, *scientist*.

## Results

We compared the mean values between both conditions, i.e. we were interested in whether there is a significant difference between the commitment to fulfill an idealized function and

the idealized function itself. An independent t-test resulted in a significant result  $t(115) = 2.164$ ;  $p = 0.003$ . The mean values are shown in Figure 3:




Figure 3: Comparison between average ratings for how surprised people would be if they encountered a person who (1) is not committed (Commitment) to carrying out her / his social role vs. (2) a person who does not fulfill the idealized function of her / his social role (Idealized Function).

## Discussion

Figure 3 compares the centrality ratings for the normative dimension (i.e. the commitment to fulfill the idealized function) and the function dimension (i.e., the function itself). Crucially, the normative dimension is significantly less mutable, hence more central, than the representation of the function. To illustrate, this means that although actually producing empirically informed theories is a central feature of the concept associated with scientist, the commitment to produce such theories is an even more central feature. In addition, the result that although not an essence, the normative dimension is a central dimension of DCCs has a very important consequence for a particular foundational challenge that could be raised against DCCs. Some theorists might accept that the normative dimension is associated with the concepts expressed by social role terms such as *scientist* and *artist*, but insist that, properly speaking, it is not part of their conceptual structure. If this was the case, the normative dimension should be highly mutable, which is inconsistent with the results of this experiment. For Study 2 shows that the normative dimension is an integral part, and possibly one of the most central parts, of the conceptual structures of certain social role concepts.

## 5. General Discussion

We have argued that the content of the normative dimension represents commitments to fulfill certain idealized functions. So the structure of DCCs includes, at least, the following dimensions: perceptual (how typical scientists look), functional (function of scientists), and normative (commitment to perform a certain function). Furthermore, the normative dimension seems to be, in the case of strong DCCs, a relatively central and immutable dimension.

Knobe et al. (2013) raise an important challenge to any

theory of DCCs. We have argued that the lexical concepts expressed by words such as *scientist* are DCCs with a normative dimension. The precise content of the normative dimension was postulated to explain the linguistic behavior of the corresponding nouns. Assuming our account explains the relevant data, we can say that it is ‘descriptively adequate’. However, why should there be a normative dimension in the conceptual structure of DCCs at all? In particular, why should social roles, but not other types of concepts, have a normative dimension? If we want our account of DCCs to also meet the demands of ‘explanatory adequacy’, we have to confront these questions. We have to explain why, given the basic function of concepts and categorization in cognition, and given the (metaphysical) properties of classes such as social roles, it should turn out that the normative is part of the conceptual structure of concepts such as *scientist* and *artist*.

Following Rosch (1999), we assume that the formation of prototypes is guided by two basic principles of categorization, one having to do with the function of a conceptual system, and the other with the metaphysical structure of the perceived world:

*Cognitive economy.* The goal of our conceptual system is to provide maximum information about the environment with the least cognitive effort. It is to an organism’s advantage to have as many properties as possible predictable from knowing any one property.

*Perceived world structure.* The perceived world is not an unstructured set of equiprobable co-occurring features. Rather, objects and events in the world are perceived as having high-correlational structure.

To see how these principles interact in concept formation, consider an example. Suppose you interact with objects  $y$  and  $z$ . How will they behave in possible future scenarios? To answer this question, you need to categorize  $y$  and  $z$ . Suppose  $y$  and  $z$  are golden yellow, but that  $y$  is static whereas  $z$  moves in all sorts of elaborate ways. At some point, you conclude that  $y$  is a gold ring and  $z$  is a bee. You can then make many predictions about their behavior, within a range that is useful for most ordinary purposes. Now, the behavior of the bee is much more complicated, and you might have to use some form of means-ends reasoning. Still, for most ordinary purposes, the repertoire of basic goals which we use to interact with a random bee—e.g., that it wants pollen and will protect the hive—is based on its kind and context. For some purposes, we might subdivide the kinds of bees, and this will determine a more refined set of functions, but even in this case there is clearly no use in representing bees in terms of the normative dimension, i.e., in terms of their ‘commitments’.

However, to add flexible social entities such as humans onto the perceived environment is to add a key layer of complexity. The behavior of humans is, given the demands of everyday life, radically unpredictable from the sorts of general biological facts we use to understand other animals. To predict human behavior in any useful way, one needs to know more than just the general biological kinds under which they fall and certain facts of the environment. In particular, one needs to know their social roles and the functions of those roles. Furthermore, one key property of

social roles is that, within certain parameters, they are often voluntary. Since Jack can choose to be a salesman or scientist, and since Jack can also choose to change roles, it is useful to categorize Jack in terms of his commitments. To generalize: human behavior is such that knowing their commitments is fundamental to make useful categorizations and predictions. In particular, it is useful to know not only whether Jack is good or bad at something, but also whether he is committed to it. Suppose you know that Jack is a really good pianist. Will he become a professional pianist? Hard to say. Suppose that you know, in addition, that Jack is committed to being a pianist. Then predicting his future behavior is easier. Suppose, on the other hand, that John is not a particularly good pianist. Will he become a professional pianist? Unlikely, but still hard to say. Suppose you find out that he is committed to being a pianist. You can predict, with some confidence, that he will keep playing, despite his current level.

Almost everyone is, at some point in their development, quite bad at performing the functions of their future social roles; however, whether they will eventually perform those functions, and whether they will become good at them, depends to a large extent on their commitments. That social role concepts encode the function and normative dimensions allows us, in various contexts, to make entirely different predictions, crucial to determine how someone will behave in certain conditions. Hence, despite the high-level talk of 'abstract values' and 'normative dimensions', there is nothing mysterious about why, given the basic function of categorization and certain basic properties of the flexibility of human behavior, many of our social role concepts should include a dimension that encodes the relevant commitments. Without this dimension, our predictions regarding the behavior of others would be substantially impoverished.

This account of the explanatory adequacy of our theory of DCCs explains why there are weak and strong DCCs. In our view, DCCs capture certain aspects of the basic structure of the perceived human social world, including sets of co-occurring features. Now it is widely believed---often correctly---that membership in some social roles requires a greater degree of commitment and effort than membership in others. *Cognitive Economy* and *Perceive World Structure* entail that social role concepts should trace these differences. If we compare the weak (e.g., *doorman*) and strong DCCs (e.g., *scientist*), we can easily see that their orderings (see Figure 1 above) reflects a common cultural valuation along those lines.

## Conclusion

We have seen that various social role concepts are dual character concepts (DCCs), and others are either weak or can be easily coerced into DCCs. We defended an account of the normative dimension of DCCs according to which this dimension represents commitments to fulfill the idealized function of a certain social role. We also showed that this dimension is central to their conceptual structure. Judgments involving social role concepts such as *artist*, *scientist* and *friend* are an essential part of social cognition, hence we expect that our account of the content of DCCs

will have implications for theories of everyday social categorization and judgments.

## References

- Gelman, S. (2003). *The essential child: Origins of essentialism in everyday thought*. New York: Oxford University Press.
- Knobe, J., Prasada, S., Newman, G.E. (2013). Dual character concepts and the normative dimension of conceptual representation. *Cognition*, 127, 242–257.
- Lakoff, G. (1973). Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2, 458-508.
- Laurence, S., Margolis, E. (1999). Concepts and cognitive science. In E. Margolis and S. Laurence (Eds.), *Concepts: Core Readings*, Chapter 1, pp. 3–81. Cambridge, MA: The MIT Press.
- Leslie, S.-J. (2014). “Hillary Clinton is the only man in the Obama administration”: Dual character concepts, generics and gender. *Analytic Philosophy*. forthcoming.
- Rosch, E. (1999). Principles of categorization. In E. Margolis and S. Laurence (Eds.), *Concepts: Core Readings*, Cambridge, MA: The MIT Press.
- Sloman, S. A., Love, B. C., Ahn, W. (1998). Feature Centrality and Conceptual Coherence. *Cognitive Science*, 22(2), 189-228.