Title

Gender Categories as Dual-Character Concepts?

Authors

Guo, GCai
Dweck, Carol S
Markman, Ellen M

Gender Categories as Dual-Character Concepts?

Cai Guo (caiguo@stanford.edu), Carol S. Dweck (dweck@stanford.edu), Ellen M. Markman (markman@stanford.edu)
Department of Psychology, Stanford University, Stanford, CA 94305 USA

Abstract

The folk theory of gender seems to involve two contradictory beliefs that people can hold simultaneously. One belief is that gender is biologically determined and immutable, and the other is that one has to earn gender membership by following gender norms or otherwise risk disqualifying oneself as a real member of the gender category. To explain this contradiction, as Leslie (2015) suggested, we turned to the dual-character concept framework proposed by Knobe, Prasada, and Newman (2013). Within this framework, we examined whether gender has two separate, parallel dimensions for evaluating category membership such that one can be a member in one sense but not the other. We found that gender concepts appeared dual-character-like in metalinguistic judgments but not in judgments of specific individuals who violate prescriptive gender norms identified by previous research. We might be witnessing a historical change where gender categories remain dual-character-like, but adherence to specific gender norms is no longer seen as definitional.

Keywords: gender, dual-character concepts, categorization, normative judgment.

Introduction

In the folk psychology of gender, there is a puzzling contradiction in how people mentally represent gender as a category. On the one hand, people are susceptible to psychological essentialism: they often believe that if you were born a man/woman, you will always be a man/woman, making gender the most essentialized social category (Prentice & Miller, 2007). On the other hand, it is also true that there have been rules that one must follow and traits one must exhibit to qualify as a "real man" or "real woman" (Leslie, 2015), or risk disqualifying oneself as a real member of the gender category (Vandello & Bosson, 2013). Thus, a person who believes that John was born a man and will always be a man, might also believe that John is not truly a man because he is highly emotional and never sticks up for himself.

How should we make sense of such a contradiction in lay beliefs about gender? Leslie (2015) argues that there might be two independent, parallel dimensions to gender categories, corresponding to two different senses of the same concept. Leslie illustrated her argument with the example of Hilary Clinton, who was referred to as the "only man in the Obama administration". Hilary Clinton is clearly *not* a man on the concrete dimension, but may be viewed as a man on the dimension of the abstract gendered values, such as being achievement-driven and having power and personal strength. In Leslie's view, the existence of two such separate, independent dimensions in the mental representation of gender makes gender categories dual-character concepts, an interesting type of concept proposed by Knobe et al. (2013). Indeed, across a series of studies,

Knobe et al. (2013) demonstrated more generally that dual-character concepts allow for two independent dimensions on which normative judgments about category membership operate, such that people could think of the individual as a member of the category in one sense but not the other. Importantly, Knobe et al. (2013) found that only categories to which we attribute abstract values and for which we form normative expectations can be candidates for dual-character concepts. For example, even though both "artist" and "bus driver" are social roles, we only form expectations based on abstract values for artists but not for bus drivers—that is, it does not sound natural to say someone is a "true" bus driver, and a bus driver who is capable of driving but does not care about driving would still be a bus driver, whereas an artist who creates art only for money and does not care about creating art that inspires people would be considered an artist in one sense but not truly an artist in another sense.

In this paper, we present the first empirical examination of whether gender categories are conceptually represented as dual-character concepts in people's lay beliefs. Specifically, we replicated all four experiments in Knobe et al. (2013), adding gender categories, to determine whether they function as dual-character concepts.

Experiment 1

Our first experiment was a direct replication of Exp.1 in Knobe et al. (2013) with the addition of two gender concepts. Specifically, Knobe et al. argued that one way of testing dual-character concepts is to see whether the concepts can be naturally described with both "good" and "true" adjectives. Knobe et al. found that when participants judged whether statements sounded natural, dual-character concepts were rated significantly higher than were other concepts when described with the "true" adjective in the statements (e.g. she's a true artist), although not when they were described with the "good" adjective (e.g. she's a good artist). Following Knobe et al., the current experiment aimed to assess whether gender concepts resemble dual-character concepts more than other concepts in the extent to which they can be described with both "true" (a true man), which concerns abstract traits, and "good" (a good man), which concerns concrete traits.

Method

Participants. We recruited 161 participants on Amazon Mechanical Turk (38.5% female; 60.9% male; 0.6% non-binary; average age: 32.48; age range: 19-68).

Materials and Procedure. Participants completed an online survey with a randomized order of 42 pairs of statements (20 dual-character, 20 control, and 2 gender concepts). All

the statements, except those for gender concepts, were directly adopted from Exp. 1 in Knobe et al. (2013). Each pair of statements for a specific concept contained a "good" statement (e.g. Marie is a good artist) and a "true" statement (e.g. Marie is a true artist) that indicated possession of concrete traits and abstract values, respectively. Participants rated both types of statements on a 7-point Likert scale, ranging from 1(*sounds weird*) to 7(*sounds natural*).

Results and Discussion

We first compared whether "man" and "woman" differed significantly from each other in participants' ratings for the two statements by testing a linear mixed-effects (LME) model with Satterthwaite approximations for degrees of freedom in R (see https://osf.io/f5y2q/ for model parameterization and all codes for analyses). No difference was revealed between "man" and "woman", as there was neither a significant interaction between concept and statement types, $F(1,320)=2.44$, $p=.12$. nor differences in planned comparisons for "good" and for "true" statements, bs <0.20, ts<1.75, ps>.30. Thus, we treated them as a single category "gender" in subsequent analyses.

We then examined whether there were significant differences among dual-character, control, and gender concepts in how natural participants perceived the "good" and "true" statements to be for each kind of concepts. The results revealed a significant interaction between concept and statement types, $F(2,39)=11.20$, $p<.001$. The interaction was driven by the fact that while the three types of concepts did not differ significantly from one another in participants' ratings for "good" statements, $F(2,39)=0.19$, $p=.83$, they differed significantly in "true" statements, $F(2,44)=27.72$, $p<.001$, such that both dual-character and gender concepts had higher ratings in "true" statements than did control concepts (bs>1.34, ps<.013), and gender concepts did not differ from dual-character concepts in this regard, b=-0.35, SE=0.51, p=0.50 (see Figure 1).

In conclusion, the current experiment fully replicated Knobe et al.'s results such that dual-character concepts did not differ from control concepts in "good" statements but had considerably higher ratings than control concepts for "true" statements. Moreover, gender concepts were indistinguishable from dual-character concepts.


Fig. 1: Means of participants' (*N*=161) ratings for "good" and "true" statements for each type of concepts in Exp.1. Error bars represent 95% confidence intervals.

Experiment 2

Knobe et al. (2013) argued that if one potential member of a dual-character category has sufficient concrete traits but lacks the abstract traits, then this person would be considered a member of that category in one sense but also not truly a member of that category. Knobe et al. (2013) demonstrated that after reading a vignette that depicted such a scenario, participants would equally endorse the statement "there is a sense in which X is a…" (member statement) and the statement "ultimately if you think about what it really means to be a …, you'd have to say X is not truly a …" non-member statement) only for dual-character concepts. Following Knobe et al. (2013), we conducted Experiment 2 to examine if participants would also endorse both kinds of statements for gender categories after reading comparable vignettes. That is, they would agree both that "there is a sense in which the character is a man" and that "ultimately if you think about what it really means to be a man, you'd have to say the character is not truly a man".

Method

Participant. We recruited 153 participants on Amazon Mechanical Turk for Experiment 2. The final sample included 147 participants, as six participants did not answer the attention check question correctly (59.2% male; 40.8% female; average age: 37.24; age range:19-77).

Materials and Procedure. We adopted the same vignettes from Knobe et al. (2013) for dual-character, natural kind (e.g. chicken), and control (e.g. cashier) concepts, and created our own vignettes for gender concepts. We used the same concepts here as in Knobe et al., which were the ten most dual-character-like and the ten most control-like concepts, rather than the full set, from their preliminary study. All the vignettes described someone or something that had sufficient concrete traits but lacked the abstract values/features associated with a certain category. Extending this paradigm to gender, we created vignettes based on research examining prescriptive norms for "man" (Vandello, Bosson, Cohen, Burnaford, & Weaver, 2008) and for "woman"(Heilman, 2001). Specifically, the "man" vignette was "John is the father of two children. However, John is very emotional, as he cries for many small things in everyday life. John is also a stay-at-home dad who has never contributed anything to the family's finance and is very submissive to his wife. Moreover, he never sticks up for himself or his family when challenged by others". The "woman" vignette was "Linda is the mother of two children. However, Linda is very dominant, as she never shows even the slightest hint of weakness and is considered intimidating by many of her colleagues in construction. Linda is the breadwinner of her family and is very bossy with her husband. Moreover, she always sticks up for herself and her family when challenged by others".

Participants received all the 32 vignettes in a randomized order, and after reading each vignette, they first rated the concrete member statement and then the "ultimately a non-

member" statement on a 7-point Likert scale ranging from 1(*false*) to 7(*true*).

Results and Discussion

As in Exp. 1, we first conducted the same LME model to compare "man" and "woman": there was a significant interaction between concept and statement type, $F(1,292)=57.56$, $p<.001$. Specifically, even though for both "man" and "woman", member statements were rated much higher than were non-member statements, the gap between ratings for the two statements was larger for "woman", $b=-5.18$, $SE=0.12$, $p<.001$, than for "man", $b=-4.00$, $SE=0.17$, $p<.001$. We therefore separated "man" and "woman" in subsequent analyses.

We then examined whether the five types of concepts differed significantly in how participants rated the two types of statements for the concepts. The results revealed a significant interaction between concept and statement types, $F(4, 27)=37.59$, $p<.001$ (see Figure 2). Planned contrasts first showed that the results in Knobe et al. (2013) successfully replicated: the difference between participants' ratings of the two types of statements differed significantly for dual-character and control concepts, $F(1,5712)=722.88$, $p<.001$ and also for dual-character and natural kind concepts, $F(1,5712)=1526.53$, $p<.001$. Specifically, for dual-character concepts, non-member statements did not differ significantly from member statements, $b=-0.14$, $SE=0.50$, $p=.79$, whereas for control concepts, non-member statements had significantly lower ratings than did member statements $b=-2.76$, $SE=0.42$, $p<.001$, and for natural kind concepts, non-member statements in fact had significantly higher ratings than did member statements $b=3.59$, $SE=0.36$, $p<.001$.

We then conducted planned contrasts to examine how gender concepts compared to other types of concepts. The results showed that both "man" and "woman" were similar to control concepts but significantly different from all other types of concepts in how the two types of statements differed from each other in participants' endorsement (see Table 1).

Table 1: F-statistics for specific comparisons between concept types (C = "control", D = "dual-character", M = "man", N = "natural kind", W = "woman") in Exp.2.

	Main Effect of Concept	Interaction (Concept X Statement)
M/D	$F(1,9)=0.38$, $p=.55$	$F(1,9)=5.90$, $p=.04$
M/C	$F(1,9)=2.72$, $p=.13$	$F(1,9)=0.91$, $p=.36$
M/N	$F(1,2939)=1.20$, $p=.27$	$F(1,9)=47.38$, $p<.001$
W/D	$F(1,9)=0.02$, $p=.89$	$F(1,9)=10..04$, $p=.01$
W/C	$F(1,9)=1.63$, $p=.23$	$F(1,9)=3.46$, $p=.10$
W/N	$F(1,2246)=149.9$, $p<.001$	$F(1,9)=63.21$, $p<.001$

We then obtained Bayes factors for a set of regression models, each of which equated one type of concepts with another one of the five types of concepts, and all other types of concepts remained unchanged. We compared all the models to examine which type of concepts would yield the greatest Bayes factors when equated with gender concepts as a single category in the model and thus determine the type of concepts that gender concepts were closest to (see Table 2). The results showed that the control concepts were the closest to both "man" and "woman", as the Bayesian models that equated "man" or "woman" with control concepts yielded the largest Bayes factors, which exceeded, by large margins, the Bayes factors from all other models that equated "man" or "woman" with natural kind or dual-character concepts.

Table 2: Bayes factors for linear mixed-effect regression models that treated "man" or "woman" as the same as one of the other three types of concepts in Exp.4.

Model	Bayes Factor
Dual-Character = Man	$2.70*10^{862} \pm 1.69\%$
Control = Man	$1.18*10^{921} \pm 1.47\%$
Natural Kind = Man	$1.57*10^{692} \pm 3\%$
Dual-Character = Woman	$2.44*10^{815} \pm 2.80\%$
Control = Woman	$1.13*10^{902} \pm 1.60\%$
Natural Kind = Woman	$4.06*10^{617} \pm 2.35\%$

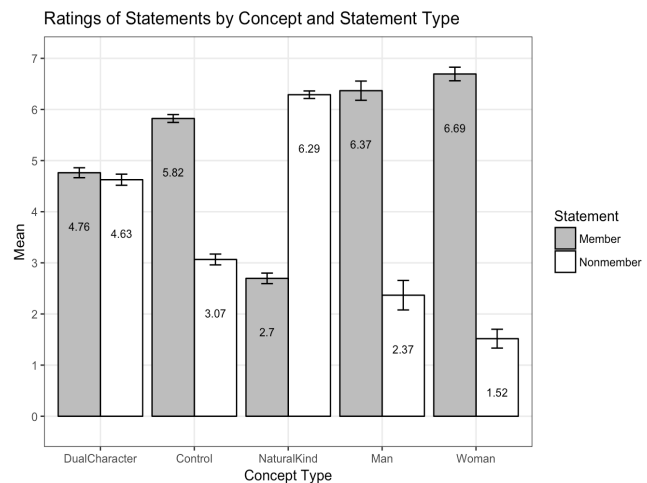Ratings of Statements by Concept and Statement Type

Fig. 2: Means of participants' (N=147) ratings for "member" and "non-member" statements for each type of concepts in Exp.2a. Error bars represent 95% confidence intervals.

In summary, the results from Experiment 2 suggested that when participants based their judgments on vignettes where the individual possessed sufficient traits on the concrete dimension but lacked core features on the abstract dimension, gender concepts did not resemble dual-character concepts as they did in Experiment 1, but rather became closer to control concepts.

Experiment 3

So far, we have found that when participants made decontextualized metalinguistic judgments (Exp. 1) gender concepts looked like dual-character concepts. They judged that saying "X is a true man/woman" is an acceptable sentence. However, in Experiment 2 when participants had to judge a particular character with traits and properties fleshed out in a vignette, gender concepts no longer appeared dual-character-like. Participants disagreed, for

example, that "there's a sense in which John is a man, but ultimately if you think about what it really means to be a man, you'd have to say John is not truly a man". We now turn to replicating Knobe et al. (2013)'s Exp.3, which returned to rating how natural two different kinds of statements sound in the absence of vignettes. Following Knobe et al., the current experiment examined whether it was natural to say someone could be a man but ultimately not a man, or someone is clearly not a man but after all could be seen as a man.

Method

Participants. We recruited 150 participants on Amazon Mechanical Turk. Three participants failed to pass the attention check question, leading to a final sample of 147 participants (49.7% female; 50.3% male; average age: 34.82; age range: 18-67).

Materials and Procedure. The same ten dual-character concepts, ten control concepts, and two gender concepts from Exp. 2 were included in the current experiment. For each specific concept, two statements were constructed: (1) "There's a sense in which X is clearly a…, but ultimately, if you think about what it really means to be a X, you'd have to say that there is a sense in which X is not a…at all" (ultimate non-member statement), and (2) "there's a sense in which X is clearly not a…, but ultimately, if you think about what it really means to be a…, you'd have to say that there is a sense in which X is a true…after all" (ultimate member statement). Participants received the statements for each concept in a randomized order, and for each concept, participants rated both statements on a 7-point Likert scale ranging from 1(*sounds weird*) to 7(*sounds natural*).

Results and Discussion

We conducted the same type of LME models as in our previous experiments to examine whether participants' ratings for the two types of statements differed for different types of concepts. The results first showed that "man" and "woman" differed significantly from each other, as overall, across statement types, "man" had higher ratings than did "woman", $F(1,146)=20.57$, $p<.001$, $b=-0.47$, $p<.001$, and there was a significant interaction between concept and statement type, $F(1,146)=4.88$, $p=.03$. Therefore, we separated "man" and "woman" in the subsequent analyses.

We then examined whether dual-character, control, and gender concepts differed significantly from one another (see Figure 4) and found that "man" was considerably similar to dual-character concepts but different from control concepts, whereas "woman" was significantly similar to control concepts but different from dual-character concepts. Specifically, the results first revealed a significant interaction between concept and statement type, $F(3,19.26)=3.15$, $p=.05$, such that the difference between the two types of statements varied for different types of concepts. To interpret the interaction, we first compared dual-character concepts to control concepts. We found that the results in Knobe et al. (2013) successfully replicated, such that there was both a significant main effect of concept,

$F(1,27.29)=66.77$, $p<.001$, driven by higher ratings for dual-character concepts than for control concepts across statement types, $b=1.17$, $SE=0.14$, $p<.001$, and a significant interaction between concept and statement type, $F(1,18.87)=6.62$, $p=.02$. Specifically, the interaction was driven by the fact that there was no difference between member and non-member statements for dual-character concepts, $b=0.02$, $SE=0.14$, $p=.90$, whereas for control concepts, member statements had significantly higher ratings than did non-member statements, $b=-0.33$, $SE=0.10$, $p=.002$. We then conducted planned contrasts using linear mixed-effect models to how gender concepts compared to dual-character and control concepts (see Table 3).

Table 3: F-statistics for specific comparisons between concept types (C = "control", D = "dual-character", M = "man", W = "woman") in Exp.3.

	Main Effect of Concept	Interaction (Concept X Statement)
M/D	$F(1,10)=1.28$, $p=.28$	$F(1,9)=0.84$, $p=.38$
M/C	$F(1,10)=12.41$, $p=.005$	$F(1,10)=14.32$, $p=.003$
W/D	$F(1,11)=6.34$, $p=.03$	$F(1,9)=0.001$, $p=.97$
W/C	$F(1,10)=3.015$, $p=.11$	$F(1,10)=3.017$, $p=.11$

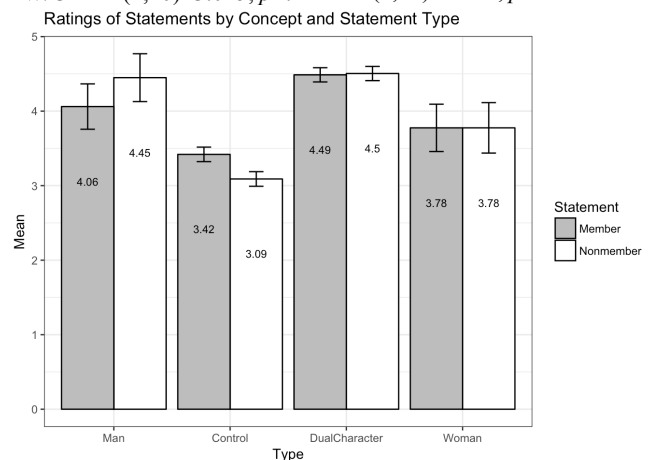Ratings of Statements by Concept and Statement Type

Fig. 4: Means of participants' (*N*=147) ratings for "member" and "non-member" statements for each type of concepts in Exp.3. Error bars represent 95% confidence intervals.

The results showed that "man" was very similar to dual-character concepts, as there was no main effect or interaction between concept and statement type. In contrast, "man" differed significantly from control concepts, such that there was both a significant main effect of concept and a significant interaction (see Table 3). The main effect was driven by the fact that overall across statement types, "man" had significantly higher ratings than did control concepts, $b=1.16$, $SE=0.33$, $p=.005$. The interaction was driven by the fact that for "man", as previously reported, the non-member statement had significantly higher ratings than did the member statement, whereas for control concepts, the non-member statement had significantly lower ratings than did the member statement.

We then compared "woman" and dual-character concepts, and the results showed that although there was no

1757

significant interaction between concept and statement types, there was a significant main effect such that "woman" overall had significantly lower ratings than did dual-character concepts, b=-0.72, SE=0.28, p=.03 (see Table 3). Similarly, we compared "woman" and control concepts, and the results showed that there was neither a significant main effect of concept nor a significant interaction between concept and statement type (see Table 3). Therefore, it seems that "woman" is closer to control concepts than to dual-character concepts. To further investigate the distance between gender concepts and dual-character and control concepts, we conducted the same type of Bayesian model comparison as in our previous experiments to compare the Bayes factors obtained (see Table 4). The model comparisons showed that "man" was much closer to dual-character concepts than to control concepts in this case. In contrast, consistent with the results from the linear-mixed effect models, "woman" was much closer to control concepts than to dual-character concepts in this case.

Therefore, Exp.3 showed that in participants' metalinguistic judgments about whether it sounded natural to describe something or someone as being a member of a category in one sense but not in the other, "man" resembled dual-character concepts, such that it made sense to say someone is a member of a category in one but not the other sense. In contrast, "woman" resembled control concepts, such that, in comparison to "man" and dual-character concepts, it made less sense to say that one is a member in one but not the other sense.

Table 4: Bayes factors for linear mixed-effect regression models that treated "man" or "woman" as the same as one of the other three types of concepts in Exp.3.

Model	Bayes Factor
Dual-Character = Man	$1.10*10^{130} \pm 1.67\%$
Control = Man	$2.77*10^{113} \pm 2.09\%$
Dual-Character = Woman	$1.22*10^{123} \pm 2.80\%$
Control = Woman	$4.82*10^{126} \pm 1.60\%$

Experiment 4

Knobe et al. worried that qualifications such as "ultimately" and "there's a sense" might have influenced participants' judgments of the vignettes in Exp.2. Therefore, our Exp. 4, a direct replication of Exp.4 in Knobe et al. (2013), tested whether the results in Exp.2 would hold up without these qualifications.

Method

Participant. We recruited 150 participants on Amazon Mechanical Turk, and seven participants failed the attention check question. The final sample therefore included 143 participants (40.6% female; 59.4% male; average age: 33.12; age range: 18-71).

Materials and Procedure. Exp. 4 used the same vignettes as in Exp. 2, but half of the participants were randomly assigned to rate only member statements (e.g. John is a

man) and the other half assigned to rate only non-member statements (e.g. John is not a man) after reading the vignettes.

Results and Discussion

We conducted the same LME models as in Exp. 2 for Exp. 4 with adjustment of random effects due to the between-subjects design for statement type in the current experiment. We first examined whether "man" and "woman" differed from each other and found a significant interaction between concept and statement type, $F(1,141)$=15.25, p<.001. We therefore separated "man" and "woman" in subsequent analyses.

We then examined how dual-character, control, natural kind, and gender concepts compared to one another (see Figure 5). The results showed a significant interaction between concept and statement type, $F(4,27.73)$=22.57, p<.001. Planned contrasts between gender concepts and all other types of concepts showed that "man" and "woman" differed significantly from all other types of concepts except for control concepts (see Table 5).

The results in Exp. 4 were therefore similar to the results in Exp. 2, where gender did not function as a dual-character concept when participants read specific vignettes that indicated the specific information of the characters.

Table 5: F-statistics for specific comparisons between concept types (C= "control", D = "dual-character", N = "natural kind", M = "man", W = "woman") in Exp.4.

|	Main Effect of Concept	Interaction (Concept X Statement)
M/D	$F(1,765)$=1.51, p=.22	$F(1,9)$=5.10, p=.05
M/C	$F(1,7.99)$=0.19, p=.67	$F(1,9)$=0.74, p=.41
M/N	$F(1,137)$=8.87, p=.003	$F(1,9)$=15.68, p=.003
W/D	$F(1,599)$=1.26, p=.26	$F(1,9)$=8.00, p=.02
W/C	$F(1,8.07)$=0.05, p=.83	$F(1,9)$=2.98, p=.12
W/N	$F(1,164)$=13.96, p<.001	$F(1,10)$=19.12, p=.002

Fig. 5: Means of participants' (N=143) ratings for "member" and "non-member" statements for each type of concepts in Exp.4. Error bars: 95% CIs.

General Discussion

We built on the studies conducted by Knobe et al. (2013) to examine whether gender categories could be represented as

dual-character concepts in folk conceptual structure. Our current findings revealed that gender concepts (especially "man") resembled dual-character concepts in decontextualized metalinguistic judgments but not when the judgments were about individual people with particular traits. That said, across experiments, man looked more like a dual-character concept than did woman.

First, why might gender concepts resemble dual-character concepts only in metalinguistic judgments? One possibility is that in Exp. 2 and 4, with specific vignettes detailing the counter-stereotypical features and explicit questions asking about the truth value of the statements, participants might have been more susceptible to a social desirability bias than were those asked to judge decontextualized metalinguistic judgments in Exp. 1 and 3. However, we have some recent findings showing that framing the statements as reflecting *society's* perspective and thus not asking about participants' own endorsement did not alter the patterns.

Another possible explanation is that in Exp. 1 and 3 which called for metalinguistic judgments, participants were acknowledging the naturalness of stating, for example, that someone "is not a true man" without having to endorse specific reasons for the disqualification. In contrast, in Exp. 2 and 4, where participants read vignettes depicting specific qualities of a character that violated stereotypical gender norms, they denied that the character was "not a true man". This suggests that participants might no longer accord definitional power to the counter-stereotypical behavior reflected in the vignettes, which were based on long standing gender norms that might now be considered historical vestiges. Specifically, recent advances in legal and social policies could have influenced people's perceptions of social norms. For instance, a recent longitudinal study on perceptions of sexual orientation suggests that the legalization of same-sex marriage in 2015 radically changed participants' perceptions of social norms regarding sexuality, such that they became much more likely to say that acceptance of homosexuality is the norm of the current American society (Tankard & Paluck, 2017). A recent linguistic study also revealed that the adjectives used to describe "man" and "woman" also changed greatly over the past century, with the change being especially prominent for "woman" (Garg, Schiebinger, Jurafsky, & Zou, 2018). Thus, people's beliefs about gender and the related stereotyped norms and roles might have changed significantly from those detailed in the vignettes. In other words, the metalinguistic judgments reflect a long-standing, historical representation of gender as a dual-character concept where it makes sense to think of someone as a true man or woman, but, if people are no longer holding rigid stereotypic gender roles, then they would reject those concrete instantiations depicted in the vignettes. Our ongoing studies will further probe this possibility.

The second question of interest that arose from the current results is the difference in judgments about "man" versus "woman". Across three of our four experiments "man" resembled dual-character concepts more than did "woman". Research on precarious manhood provides helpful explanations for such a pattern: it is well-documented that across cultures, manhood is viewed as something that is earned, needs to be proved, and can be easily lost, whereas womanhood is considered a natural product of biological maturation, and thus it is not so easily lost once it is earned (Vandello & Bosson, 2013). Such a difference in the social construction of manhood and womanhood across cultures is likely the reason why "man" appeared more dual-character-concept-like than did "woman" in the current study.

In conclusion, the current study showed that gender categories resembled dual-character concepts in meta-linguistic contexts but not in substantiated contexts with detailed portrayals of the traits on each dimension. An intriguing possibility is that we are witnessing an historical change in how people view broad gender stereotypes. The intuition that gender is a dual-character concept where it seems natural and sensible to judge that someone is "not a true man" remains robust, but the link between that intuition and particular stereotyped beliefs may be eroding.

References

Garg, N., Schiebinger, L., Jurafsky, D., & Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, *115*(16), E3635-E3644.

Heilman, M. E. (2001). Description and prescription: How gender stereotypes prevent women's ascent up the organizational ladder. *Journal of Social Issues, 57*(4), 657-674.

Leslie, S. (2015). "Hillary Clinton is the only man in the Obama administration": Dual character concepts, generics, and gender. *Analytic Philosophy,56*(2), 111-141.

Knobe, J., Prasada, S., & Newman, G. E. (2013). Dual character concepts and the normative dimension of conceptual representation. *Cognition,127*(2), 242-257.

Prentice, D. A., & Miller, D. T. (2007). Psychological essentialism of human categories. *Current Directions in Psychological Science,16*(4), 202-206.

Rouder, J. N., & Morey, R. D. (2013) Default bayes factors for model selection in regression, *Multivariate Behavioral Research*, 47, 877-903.

Tankard, M. E., & Paluck, E. L. (2017). The effect of a supreme court decision regarding gay marriage on social norms and personal attitudes. *Psychological Science,28*(9), 1334-1344.

Vandello, J. A., & Bosson, J. K. (2013). Hard won and easily lost: A review and synthesis of theory and research on precarious manhood. *Psychology of Men & Masculinity,14*(2), 101-113.

Vandello, J. A., Bosson, J. K., Cohen, D., Burnaford, R. M., & Weaver, J. R. (2008). Precarious manhood. *Journal of Personality and Social Psychology*, *95*(6), 1325-1339.