

Durham E-Theses

An Investigation into the perception of social media endorsement posts in accordance with the attributes of the endorser, utilising a chain of cultural transmission reminiscent of the Instagram algorithm.

SUZY-MAY LISTER

How to cite:

LISTER, SUZY-MAY (2023) An Investigation into the perception of social media endorsement posts in accordance with the attributes of the endorser, utilising a chain of cultural transmission reminiscent of the Instagram algorithm. Masters thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/15329/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

An Investigation into the perception of social media endorsement posts in accordance with the attributes of the endorser, utilising a chain of cultural transmission reminiscent of the Instagram algorithm.

By Suzy-May Lister

Contents

Abstract	3
Introduction	4
Literature review	4
Methodology	10
Analysis	19
Results and discussion	29
Conclusion	32
References	34

Abstract

The following study is an investigation into social media product endorsements and how the individual promoting a product impacts how likely the post is to receive engagement from consumers. An experiment was created featuring simulations of influencer marketing posts one might see on Instagram, including a product image and profile picture. Some of the profile pictures contained variation in attractiveness, some posts contained a follower number, and some contained an occupation, or a combination of these attributes. The experiment was run via Amazon Mechanical Turk and a chain of cultural transmission was created as the first group of participants responses were collected and used to determine how frequently certain attributes would appear in the next group of social media posts shown in the experiment. The attributes that were liked most frequently in the previous set of posts were then expressed at a higher frequency for the next group of participants. This transmission chain element is unique to this study and aims to investigate whether there is a preference for attributes that are more common over attributes that are rare, while replicating the effect of social media algorithms that promote posts we have not seen before but were popular amongst others. The study produced complex results suggesting that there is no consistent effect of the frequency of representation of certain attributes on the chances of a post being liked. There was a significant preference for higher attractiveness profile pictures across all conditions and a preference for higher popularity when attractiveness was controlled. Additionally, there was a greater like frequency for relevant experts when presented alongside popularity but no significant preference for either relevant or irrelevant experts when presented alongside variation in attractiveness. Finally, there was no consistent effect of participant self-esteem on attribute like frequencies.

Introduction

While cultural variation is common cumulative culture is rare (Boyd and Richerson, 1996, p.2). Cumulative culture refers to the aggregation of traits over generations to produce behaviours that an individual could not learn purely themselves (Ibid.). The capacity to form this complex patchwork of knowledge is something unique to humans, possibly chimpanzees, and some species of bird. This suggests that it requires more than just intelligence and sociality to be apparent within a species for cumulative culture to evolve as species such as Cebus monkeys, that are among the cleverest animals alive, have not been seen to exhibit cumulative cultural evolution (Boyd and Richerson, 1996, p.9). Boyd and Richerson posit that cumulative cultural evolution is driven by observational learning, the capacity for which sets us aside from other animals (Ibid. p.6).

Cumulative culture is achieved through social learning. In almost every realm of human experience knowledge has advanced far beyond the learning capacity of a single individual, without such knowledge society today would be a very different place. As a result, social learning mechanisms are essential in order for accumulated knowledge to be passed on as learning and devising such a vast amount of knowledge yourself would be impossible. This raises the question of which social learning strategies allow certain types of cultural transmission. I would like to explore social media marketing as the context within which to examine contemporary social learning mechanisms and cultural transmission, within the realm of influencer marketing specifically.

Influencer marketing strategies employ a social media spokesperson to advertise a specific product to their followers. With the rise of social media in the past decade, followed by the COVID-19 pandemic social media use is booming. The number of social media users increased by 13.1% in the year leading up to July 2021, with 4.48 billion active users at this time (Influencermarketinghub, 2021). Moreover, '63% of marketeers intend to increase their influencer marketing budget in the next year' and businesses estimated to be making '\$5.20 for every \$1 spent on influencer marketing' according to The Influencer Marketing Hub (Santora, 2021). Therefore, the efficacy of influencer marketing strategies in increasing product sales is evident, however, there is little research questioning why this is the case – for what reason do people choose to socially learn from influencers? Are social media acting as a means of cultural transmission themselves? I would like to address these questions by analysing primary data on consumer behaviour in accordance with anthropological theories and concepts relating to cultural transmission. Through this I hope to illuminate whether forms of prestige bias play a part in the success of influencer marketing. The indicators of prestige adopted will be attractiveness, popularity and expertise.

Literature Review

In 2001 Henrich and Gil-White coined the term 'prestige bias' to describe the way in which animals may choose from whom to socially learn. This kind of bias is at play when an individual chooses to learn from another individual whom they perceive to be prestigious. There are many factors affecting who is perceived as prestigious and by whom, often pertaining to skill or knowledge. For example, you might be more likely to ask the most muscular person in the gym for advice on how to gain strength as that person we would assume has plentiful knowledge in this area. In this way, prestige bias can be adopted as an adaptive strategy that enables us to socially learn information from one another that would otherwise take a long time to learn purely ourselves. However, someone's skill level is not always immediately obvious so we might instead observe the levels of deference expressed towards individuals

by others as a good indicator of perceived prestige. Though deference itself can also be difficult to identify. However, in the digital age it is easier to identify and quantify esteem as people express their preferences online on a daily basis. Social media platforms encourage users to 'react', 'like' and 'share' content meaning there is always some gauge of the appeal of a person or a post. In this study we assume that the act of 'liking' a post on social media is exhibiting a positive preference toward that post.

You could argue that social media influencers have an assumed level of prestige as an implication of their following. Generally, the greater the number of followers the greater the conferred prestige as following somebody on social media could be considered an act of deference, especially when the follower and followed do not know each other. For this reason, popularity is an attribute under investigation as a potential source of prestige bias and hence the success of influencer marketing.

Experimental examples of popularity bias

Experimental examples of prestige-bias occurring among modern humans highlight its plausibility as a social mechanism at play in consumer behaviour. The idea of prestige as related to attention was first introduced by Michael Chance who suggested that humans socially learn from, and provide attention to, those of high status (1976, in Barkow et al., 2012). Therefore, in the context of social media, celebrities and influencers could be considered 'high status' as they have the most followers and receive the most attention. However, research on this topic is limited. Atkisson et al. conducted the first experimental test of prestige bias in adults with their study in which adults had to design arrowheads optimal for a certain type of 'hunt' (2012). Participants' designs were scored after each task meaning each individual could recognise how well they had done. Players could then opt to copy another's arrowhead for the following 'hunts' and were able to look at the arrowheads others had designed (though in reality they were predetermined) and see how long the other players had spent looking at each (score for each head was not shown here). It was highlighted to the players that the optimal arrowhead may change and so when assessing this information they should be thinking of it as choosing to copy the creator rather than the arrowhead itself. The study concluded that participants 'preferentially learned from prestigious models', demonstrating the human learning bias toward those that receive the highest levels of deference (Ibid.). Furthermore, Chudek et al. discovered that three- and four-year-old children are twice as likely to learn from adults who had been preferentially paid attention to for the ten seconds prior, than those adults that were ignored (2011). Furthermore, Little et al. tested the incidence of prestige bias in the context of mate choice (2015). Women were shown computer generated images of men and asked to rate their attractiveness. Each man was presented in combination with an image of a woman and her 'popularity' score, assuming popularity confers prestige. Results showed that women found men associated with 'popular' women more attractive than the men associated with less popular women. This suggests that the women assessing the pictures are socially learning from the 'popular' women by copying their apparent 'taste' in men. Although the relationship between the man and the woman photographed was not specified, female participants appear to have let the mere association effect their opinions of the men's attractiveness.

The above experiments exhibit human prestige-biased social learning and show deference and popularity to be indicators of prestige. Participants did not choose from whom to learn randomly, rather they chose those that appeared to be popular or were given the most attention. This concept can be carried forward into the context of social media as those with large followings naturally appear to be 'popular' and receive a lot of attention due to the number of people that view and engage with their

content. This association is corroborated by De Veirman et al. who found that Instagram influencers with large followings are seen as more likeable partially due to their assumed popularity (2017).

Experimental examples of attractiveness bias

Prestige bias with relation to attractiveness has been thoroughly researched in the social sciences. Previous literature shows bias toward attractive individuals to be common in many different contexts. For example, Klebl et al. found that individuals considered to be attractive were perceived to possess more moral traits versus those considered unattractive (2022). The experiment involved participants being shown 6 images of either attractive or unattractive individuals alongside moral and non-moral traits. They were then asked to rate how likely they thought the person in the image shown was to possess these traits. The results showed a strong bias towards attractive individuals when assuming the possession of moral traits such as fairness, and an overall slight preference for attractive people when assuming any kind of positive trait (Ibid.). The results draw upon the 'beauty is good' stereotype which is considered a halo-effect, meaning cognitive 'reasoning in which an impression formed from a single trait or characteristic is allowed to influence multiple judgments or ratings of unrelated factors' (Neugaard, no date). The images used in the Klebl et al. experiment were taken from the Chicago Face Database using the highest and lowest scoring individuals to form the 'attractive' and 'unattractive' conditions (2022). I have utilised the same process in my study, using the Chicago Face Database to form the more attractive, less attractive, and average attractiveness conditions. The results of the Klebl et al. experiment are corroborated by that of Miller. Miller also investigated the relationship between attractiveness and the attribution of positive personality traits (Miller, 1970, cited in Baker and Churchill, 2018). Miller found that on 15 out of 17 bipolar adjective scales relating to personality traits, moderately attractive and attractive individuals were associated with the positive expression of the trait and unattractive individuals with the negative (Baker and Churchill, 2018). The findings suggested that physically attractive individuals are assumed to act with greater autonomy and to be more in control of their lives than unattractive individuals who are perceived to be more likely to be subject to coercion or pressure from external factors (Ibid.). These findings are widely corroborated and may contribute to as to why attractive models are used in advertising campaigns, as the traits they are associated with may lead the consumer to believe they are more credible. Baker and Churchill suggest that the consumer may consider the advertisement 'less as a persuasive message and more as communication of informational content from a trustworthy source' (2018).

Attractiveness bias has also been a long-standing topic of interest to psychologists and social scientists in the advertising realm. For example, Baker and Churchill found that attractive models had a positive impact on the consumer's impression of print ads (Baker and Churchill, 2018). Moreover, attractiveness bias was identified in the realm of influencer marketing in a study by Koay et al. who discovered through a questionnaire that the attractiveness of Instagram influencers is positively associated with the consumers likelihood of making impulse purchases of products advertised (2021).

What we view as prestigious is subjective and dependent

Barkow recognises the human capacity for cognitive distortion; the ability to mentally distort incoming information, or to manipulate the parameters by which it is assessed and hence how it is understood (556). For example, there has been significant movement on the social media platform Instagram

regarding attitudes around body image, with many influencers campaigning to end 'fake' or 'photoshopped' content that perpetuates unrealistic body ideals. In doing this, campaigners are attempting to utilise the human ability for cognitive distortion in order to change the way bodies are perceived through altering the criteria by which we judge them. For example, by praising rolls of fat and stretch marks rather than slim waists and 'thigh gaps'. However, although cognitive distortion can be a method through which prestige and self-esteem can be acquired, Barkow notes that this method is not as effective at maintaining self-esteem as achieving prestige without distortion (Ibid.). Moreover, he asserts that those using 'excessive' cognitive distortion are less likely to adapt and be successful (Ibid.). As a result, it is no surprise that influencers are so popular because it is very difficult to change the entire culture determining the evaluation criteria for prestige, and so it is more efficient for an individual to seek prestige through changing themselves and copying prestigious others rather than changing theirs and others' perspectives.

Expertise-related bias

Prestige bias is not to be confused with skill bias as often the prestigious individual one chooses to learn from does not possess any skill in the area you are hoping to learn about. Prestige bias is shown to be exhibited only in the absence of enough information for skill bias to be adopted (Brand et al., 2021). However, Jimenez and Mesoudi assert that the usage of prestige-biased social learning strategies depends on the stability of the social and ecological environment (2019). When there is rapid social change the skills and knowledge that were once relevant quickly become outdated and people must learn and adapt accordingly. Hence, a dynamic environment necessitates efficient information transfer and so social learning strategies such as prestige bias become useful (Ibid.). Nowadays, trends in consumption change quickly due to rapid cultural transmission via social media and the internet, meaning consumers may be turning to influencers for product recommendations regardless. Part of the experiment discussed in this thesis involves fake social media posts by people whose occupations are specified either as an 'influencer', an expert relevant to the product being displayed or an expert irrelevant to the product being displayed. This will test the theory of Brand et al that prestige bias is only adopted in the absence of accurate skill information, in which case we would expect there to be bias toward attractiveness and popularity only in the conditions that do not include expertise (Brand et al., 2019).

Moreover, studies focussing on 'product-endorser fit' in advertising have shown product endorsement by domain specific endorsers to be more effective than non-domain specific endorsers. Product-endorser fit refers to the relevance of the product being advertised to the endorser's interests or expertise. For example, Kamins and Gupta's study concluded that congruence between the product and the endorser increases the believability of an endorser while evoking a 'more favourable' product attitude (1994, cited in Schouten et al., 2020). Further studies confirm that product endorsements are more effective in influencing purchase intention when the product seems to be relevant to the person advertising it (Till and Busler 2000; Fink et al. 2004).

Cultural transmission

Barkow et al. posit that social media are disrupting the evolved mechanisms of cultural transmission and explain that if Chance et al., were correct in their assertions then the mode of cultural transmission may have changed as younger people are more inclined to learn from social media personalities rather

than their parents or those in their community as they are perceived to be of higher status (Barkow *et al.*, 2012). This corroborates the idea that an influencer's audience may be using them as a genuine information source and socially learning from them. Furthermore, what information is transmitted via social media is affected by algorithms. The Instagram algorithm this study aims to resemble is that specific to the 'explore page'. The Instagram explore page is a feature enabling a user to view content from creators that they do not follow, in the hope that users will discover new content and people they would like to engage with, hence keeping them interested in the app. Multiple factors determine which posts one might be shown on the explore page including, the speed of engagement – how quickly that post was liked, shared, commented on and saved when it was posted, and how much engagement that creator has seen over the past few weeks (Das, 2021). In other words, what we see on the explore page depends largely upon what others like. Thus, Instagram analyses which content is popular and shows this content to those that have not yet seen it rather than showing content chronologically. This study seeks to explore online cultural transmission by echoing the way that Instagram boosts popular content by increasing the frequency of posts with popular attributes. After every ten people respond to the experiment the posts will be changed and the attributes associated with the most popular posts will feature more in the next set of posts and unpopular attributes will feature less. This aims to investigate whether the high representation of certain attributes leads the respondent to be more likely to engage with that content.

Barkow et al. also note however, that the incidence of social learning from social media personalities may depend upon age. Barkow cites De Backer et al. who found that while young people use celebrities as a source of adaptive information, older people tend to only have parasocial relationships with them. However, this does not mean to say that younger people do not also have parasocial relationships with influencers or celebrities online. Moreover, Little et al.'s experiment exploring prestige bias in the context of mate choice, also found that older participants were much less likely to be influenced by the popularity of the woman a male was paired with, when rating their attractiveness (2015). This corroborates the idea that younger people may be more inclined to socially learn from celebrities or influencers than older people. As a result, age was included in the demographic questions section of this study so that age could be investigated further if required.

Prestige, self-esteem and social media influencers

Barkow considers prestige to be the means through which humans maintain 'self-esteem', and that this is the result of the evolution of primate social dominance (Barkow *et al.*, 1975, p.556). Moreover, Barkow asserts that the evolution of self-esteem has occurred due to natural selection acting on the cultural environment, and the way by which humans aim to acquire prestige depends upon their social environment (Barkow *et al.*, 1975, p.555). This means there that is no one way of gaining prestige, rather it depends on a person's community and potential for social relationships. The social environment under consideration in this experiment is that of any person with sufficient access to the internet to have an Amazon Mechanical Turk account; the platform through which the survey was conducted.

According to Barkow's assertions we may view prestige as a way of acquiring self-esteem. Therefore, in the context of consumption, an individual that decides to make a purchase after viewing an influencer's advertisement may be aiming to achieve a form of prestige provided by the influencer and determined by the values that they perpetuate. For example, a skincare influencer likely perpetuates the ideal of having nice-looking skin, hence their followers buy the products they recommend in order to achieve

this kind of skincare prestige. This prestige then boosts the consumer's self-esteem. Moreover, Lee et al., found a relationship between the way that social media influencers are perceived and self-esteem. It was found that there was a positive indirect relationship between feelings of envy and affective well-being by means of the influencer acting as a source of 'inspiration', and this relationship was stronger amongst those with higher self-esteem (Lee *et al.*, 2022). Furthermore, studies show that consumers with low self-esteem are greater influenced by others (Bearden *et al.*, 1989, cited in Djafarova and Rushworth, 2017) and may seek assistance when making purchasing decisions because of this low self-confidence (Bearden *et al.*, 2001).

As a result of this possible relationship between prestige bias and self-esteem this paper has utilised the Rosenberg self-esteem scale so that the way the respondents react to the experiment can be studied in conjunction with an estimation of their self-esteem. It is likely that consumers choose to socially learn from influencers in the hope of replicating some appealing aspect of the influencer's life for themselves, therefore if one is satisfied with their life already they may be less inclined to engage with or be swayed by influencer content. As a result, the Rosenberg self-esteem scale can be used to identify whether participants with higher averaged self-esteem show less of a preference towards prestigious influencers than those with lower self-esteem. However, this experiment only explores which type of influencer a person would be more likely to engage with if they had to, it does not identify people that would ordinarily not engage with influencers at all, meaning this could be an interesting avenue for further research.

Other potential mediators of the efficacy of influencer marketing

Another theory as to why influencer marketing strategies are successful pertains to the formation of parasocial relationships that result from the type of engagement the public has with online personalities. Barkow describes the different kinds of prestige-conferring 'strategies' that are dependent upon the presence or absence of givers and receivers of prestige (556). This highlights the possibility for prestige to be conferred to an individual even in the absence of the prestige giver. This is occurring when people engage with the content of social media influencers online. A parasocial relationship is a one-way relationship in which an individual invests time, emotional energy and interest in a person who does not know that they exist (Rubin *et al.*, 1985). This is the classic kind of scenario exhibited by the lay social media user as they feel like they know the celebrities they follow because they know so much about them through their social media posts, however the celebrity does not know anything about this individual. Though the lay person is aware that themselves and said celebrity are not actually friends in reality, they still feel connected to them whether this be conscious or not. According to Reinikainen et al., it is this parasocial relationship that leads to influencer credibility as their audience perceive them to be trustworthy as they feel that they know them (2020). Therefore, the influencers that best facilitate parasocial relationships with their followers are likely to have the most credibility and be the most successful at marketing products. Moreover, studies have shown influencers to be more impactful on consumer behaviour than celebrities due to their greater credibility and relatability (Djafarova and Rushworth, 2017).

There have been many other factors investigated as potential mediators of the efficacy of an influencer or celebrity advertisement. For example, 'identification' and 'credibility' have been investigated as potential mediators (Schouten et al., 2020). Identification has also been referred to as 'attitude homophily' or 'similarity' and is used to describe the way that having a sense of commonality with an influencer or celebrity makes a person more inclined to follow their advice (Sokolova & Kefi, 2019).

Additionally, Reinikainen et al. found that audience interaction with an influencer's profile, and credibility, are correlated (2020). Commenting on someone's online content facilitates parasocial relationships with the creator when done oneself, and also appears to enhance credibility when witnessed by others; merely reading other people's comments is effective in enhancing influencer credibility. Moreover, parasocial relationships with an influencer are also shown to sometimes translate into 'brand trust' which is positively correlated with 'purchase intention' (Ibid.).

Methodology

Study overview

The experiment tests preferences for different combinations of influencer attributes; attractiveness, popularity and expertise. These attributes were presented to participants in the form of fake social media posts as well as survey questions relating to demographics and influencer perception. Participants were invited to choose 3 of the posts which they would be most likely to engage with had they come across them on their personal social media. The experiment acted like a transmission chain as the fake social media posts were changed after every ten participants' responses. This was to enable the attributes chosen the most by participants to be shown at a greater frequency for the next ten participants. The transmission chain format was chosen to replicate the Instagram algorithm that determines what we are shown on the 'explore page' as a result of which posts are popular amongst other users. Each group of posts were changed three times meaning four variations of each condition were created. The experiment was followed by two questionnaire sections; one relating to self-esteem and the other to influencer perception. Full ethical approval was obtained for this investigation via the Durham University anthropology department's ethics committee.

Hypotheses:

1. Posts are more likely to be liked if they contain attributes that are common amongst other posts.
2. Posts by influencers displaying higher attractiveness and popularity will be liked more frequently relative to lower attractiveness and popularity influencers.
3. There will be no significant preference for attractiveness or popularity where information on expertise is available.
4. Relevant experts get more clicks relative to irrelevant experts.
5. Those with high self-esteem are less affected by prestigious traits.

Experiment design

The experiment consisted of four conditions testing different combinations of influencer attributes; attractiveness, popularity and expertise. Attractiveness and popularity were treated as binary variables with higher and lower attractiveness and higher and lower popularity. Expertise had three variables; 'relevant expert', 'irrelevant expert' or 'influencer'. The four conditions being tested were as follows; attractiveness with popularity, attractiveness with expertise, popularity alone, and popularity with expertise. The conditions were chosen in order to explore the hypotheses in a way that could be most

easily interpreted. Hypothesis 2 investigates the like frequencies of higher attractiveness and higher popularity posts versus lower attractiveness and lower popularity posts. Therefore, a condition comprising higher and lower attractiveness and popularity posts was created so that the frequency of likes for these attributes could be compared directly. Additionally, popularity was tested individually in its own condition, as an Instagram profile cannot exist without featuring a follower number, however some accounts do not feature a person's appearance or occupation, therefore it would not be representative to test attractiveness or expertise as attributes alone without another attribute present. The results from the popularity only condition could then be compared to the results of the popularity with attractiveness condition to further understand the impact of the addition of the attractiveness attribute. It must be noted that the attractiveness with expertise condition is not representative of Instagram, however the condition is useful in order to test hypothesis 3 more broadly allowing the results to be compared to those from similar studies discussed in the literature review. The popularity with expertise, and attractiveness with expertise conditions also test hypothesis 4. This data could then be compared to the aforementioned conditions to obtain like frequencies for when the expertise attribute was not present. In doing this, sufficient data could be obtained from only four conditions while maintaining a sample size reflective of those used in similar studies in the literature review, and also staying within the Amazon Mechanical Turk budget set aside for the experiment. Each participant was tested only in one condition, and all attributes were evenly distributed in the first generation to ensure that high attractiveness was associated with higher popularity just as often as lower popularity for example.

The conditions were presented in the form of 24 fake social media posts which were displayed as a series of pictures aiming to be reminiscent of suggested posts the Instagram explore page (see Figure 1). Each post contained a picture of a product below a profile picture and a job description or number of followers, depending on the condition (see sub-section 'product genres' below for details on the images used in the posts). The font and format of the 'posts' were also chosen carefully on Canva in order to emulate Instagram via; the size and location of the profile picture, the colour and style of the circle around the profile picture, the font of the name, and the size of the product picture.



Jenni
2.9M followers



Al
550 followers
Chef





Lindsay
Tech specialist



Jack
1.2M followers



Jim
Music producer



Mel
490 followers



Claire
350 followers
Music producer



Daisy
2.5M followers
Doctor



Figure I – A selection of the posts used featuring attractiveness, popularity, and expertise attributes

The different variations of posts for each condition will be referred to as ‘generations’ alongside a number to indicate the chronology, for example, generation 3 indicates the third version of posts for a condition meaning the posts have been adapted 2 times. The first generation of images for each condition were seeded with equal proportions of attributes to ensure that the frequency of representation did not impact participants responses before the transmission chain had begun. Each starting grid was presented to 10 different participants who were instructed to choose 3 of the 24 posts presented to them in the grid to ‘like’ like they would on social media: ‘Please select 3 (no more no less) of the following posts that you would be most likely to engage with (e.g through liking, sharing, or buying the product on display) imagining you came across them on your own personal social media. Be sure to have looked through all posts before making your selection.’ 3 likes were requested however a leeway of plus or minus 1 (2 and 4 choices were accepted) was given because so many participants did not choose 3. Participants that choose more than 4 posts or only 1 post would be removed from the dataset and a new participant found to replace them to ensure 10 usable responses were acquired for each version of each grid.

After ten people had selected their images, the grid would then be changed in accordance with the proportion of likes a certain attribute received. Hence, the next version that condition would be shown to the next ten respondents displaying different proportions of certain attributes, determined by the choices of the last ten participants. The participants were not aware of this process at the time of answering. For example, for the popularity only condition, the proportion of likes for the ‘higher’ variation of popularity would be calculated and would determine the proportion of the 24 posts in the next grid displaying the higher popularity so that the next 10 participants would receive a different grid to the previous. For example, if 75% of posts chosen were higher popularity then 75% of the posts shown in the next variation of the grid would contain higher popularity and the other 25% would be low popularity.

If the proportion of likes for a certain attribute did not give rise to a whole number of posts the number of posts would be rounded to the nearest integer. For example, say that 7 out of 30 likes were for posts expressing higher popularity in grid 4; $7/30 = 0.23333333$, $0.23333333 * 24 = 5.6$, $5.6 = 6$ posts expressing higher popularity in the next version of that grid. On 2 occasions this meant that the number of calculated posts for each attribute did not add up to 24 in which case the number that was most intermediate would round whichever way were necessary to achieve 24 posts. For the grids where there were multiple attributes the attributes were treated as conjoined, for example for grid 4 the number of likes for popular influencers, unpopular influencers, popular relevant experts, unpopular relevant experts, popular irrelevant experts and unpopular irrelevant experts were counted so as not to ignore the interaction effect of the variables. Each grid was changed 3 times so that there were 4 generations in the transmission chain. Therefore, a 4 stage transmission chain was produced per condition. This enabled the proportion of attributes present in posts to be tracked over time to see if cultural transmission was occurring.

Transmission chain example – Popularity only condition

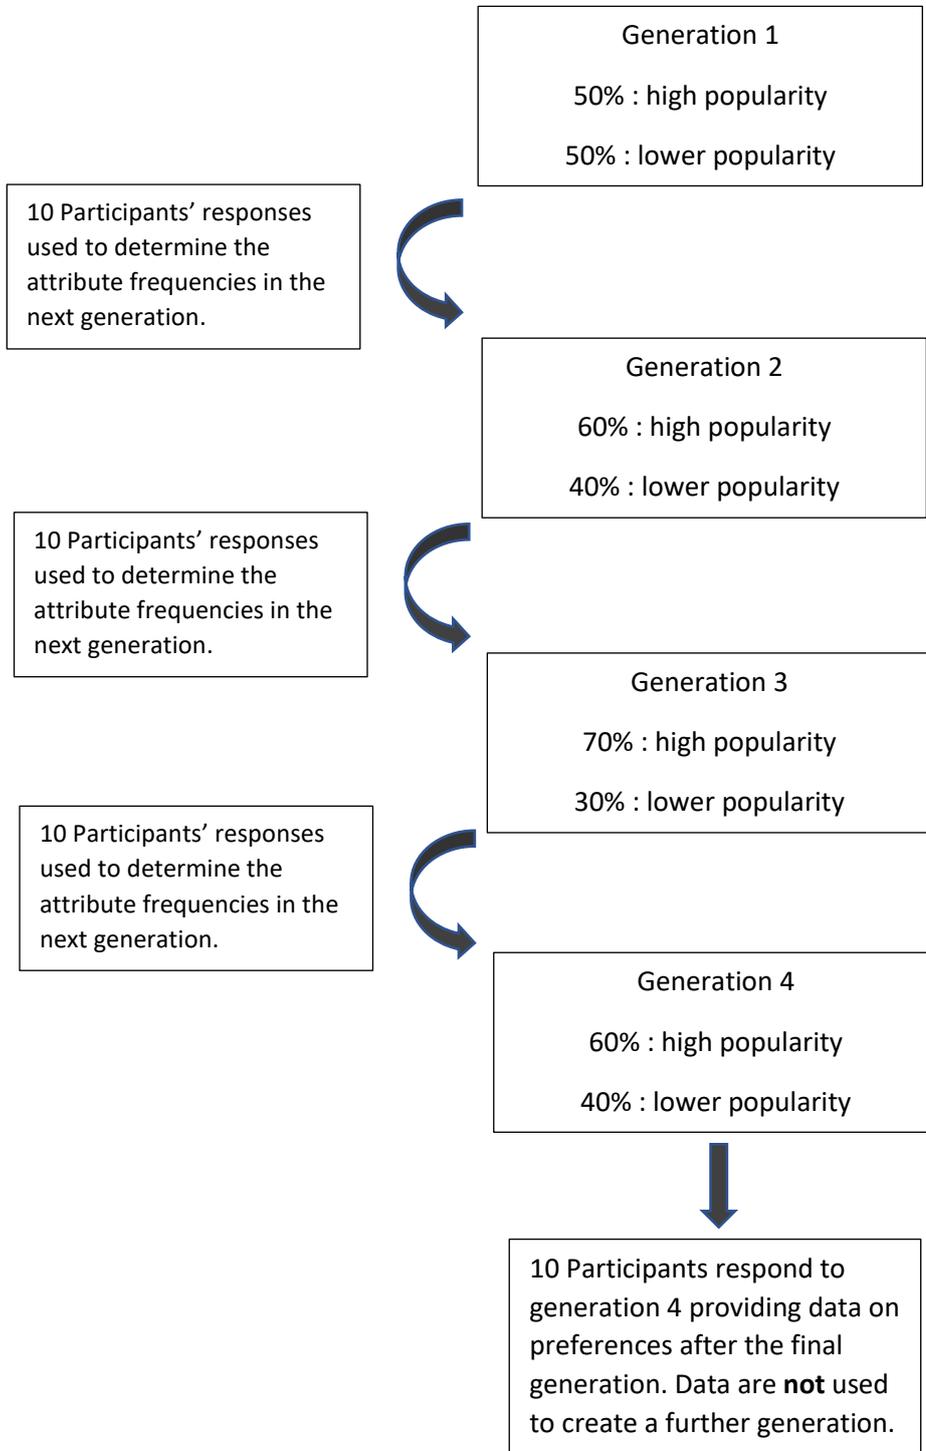


Figure 2: Example of transmission train mechanism for the popularity only condition. Values expressed are fictional.

The 'no more no less' clause in the instructions was added after the first round of responses when it was realised that a large quantity of respondents were not correctly following the instructions. It was important that each participant provided a similar number of responses to avoid the data being skewed toward those that chose to like a larger number of posts. Hence the 'no more no less' served to enhance the quality of the data while being highly unlikely to affect the results of the study. The data of those participants whose responses were not usable were deleted however the participant was still paid the same to compensate for their time.

Product genres

Instagram is an image based social media service hence an image was required to create each social media post. A broad selection of products were chosen to be featured in the posts. 24 posts were created for each generation of each combination to allow there to be four different posts across 6 product types. This was important to ensure there were sufficient product types being presented to the participant while still having multiple variations of posts for each product. The products featured on the posts were chosen as they represent the leading consumption industries; food/drink, technology, health/lifestyle and entertainment. All the products chosen have mass appeal and are not tied to specific demographics to ensure as far as possible that the participant will be viewing posts containing products that they are likely to use or consume themselves. Each type of product was presented four times with different attributes (if possible) to mitigate bias towards certain types of product. For example, in grid 3 pizza was shown both with a higher popularity person and a less popular person twice each in the first variation. The images used to create the posts were taken from Canva Pro the premium version of Canva the online graphic design platform. All images were taken from Canva Pro not a mixture of standard Canva images and Pro images to ensure the quality of the pictures used in the grids were consistent hence mitigating bias towards superior quality pictures.

The products featured were as follows:

Pizza
Coffee
Mobile phone
Water bottle
Smoothie/ juice
Camera
Laptop
Burger
Gym equipment
Vitamins/ supplements
Headphones
Journal
Ear buds

Figure 3

The attractiveness attribute

The attractiveness attribute was represented by the profile picture of the experiment's social media posts. Attractiveness was treated as a binary variable containing only a 'higher' or 'lower' variation unless it was being controlled. The profile pictures for the fake social media posts were created using the Chicago Face Database. The Chicago Face Database 'provides high-resolution, standardized photographs of male and female faces of varying ethnicity between the ages of 17-65' from the United States, intended for use in scientific research (Ma et al., 2015). The data also contains 'subjective

ratings by independent judges' of attributes such as attractiveness (Ibid.). The ratings were presented as a singular mean score for each face, acting as a proxy for perceptions of attractiveness of faces similar to those you might find on popular Western social media platforms. The 24 faces scoring the highest and lowest for mean attractiveness from the Chicago Face Database version 3.0 was selected to represent highly attractive profile pictures and lower attractiveness profile pictures. Klebl et al. chose in their study to additionally group these individuals by ethnicity whereas I did not because ethnicity of the respondent was gauged in the initial demographic survey questions hence the relationship between ethnicity of the participant and ethnicity of the person in the image shown could be investigated afterwards (Klebl *et al.*, 2022b).

The variations of attractiveness are referred to as 'higher' and 'lower' because the experiment is not concerned with categorising people as very attractive or unattractive, as this is subjective so we could not make any firm conclusions about 'highly attractive' people as we do not know that those that answered this experiment perceived them as such. This is because the people answering to this study are not the same as those that took part in the Chicago Face Database study. Hence, higher and lower variations are being used to express that on average one group of people is rated higher than the other group, we cannot definitively categorise these people as being of high or low attractiveness we only know that it is generally accepted that one group is more attractive than the other.

In the popularity only, and popularity with expertise conditions grids attractiveness was a variable to be controlled. Therefore, the median 24 faces were used to represent influencers of average attractiveness. The expressions of the faces were controlled to avoid bias towards certain facial expressions. This meant that if one of the 24 images selected to represent a certain level of attractiveness had a different facial expression to the rest, which were neutral, this face would not be used and the next best face would. For example, by incorporating the 25th most attractive face to represent the high attractiveness variable, instead of the 24th when the 24th face did not have a neutral expression.

The race of the faces used in profile pictures was not controlled as it was difficult to find sufficient very high or low scoring individuals from a single ethnic group, however all faces were taken from the United States edition of the Chicago Face Database so were considered US citizens. Hence, the highest and lowest scoring individuals were used irrespective of race to maintain a small score range within the 'higher' and 'lower' variations of attractiveness. This ensured that the difference between higher and lower attractiveness was as distinct as possible.

Condition	Mean Attractiveness Score – Chicago Face Database
Higher attractiveness	4.82-5.48
Lower attractiveness	1.52-1.89
Control/ average attractiveness	3.08-3.17

Figure 4

The expertise attribute

There were 3 types of expertise; relevant expert, irrelevant expert and influencer. The relevance of the expert relates to the type of product exhibited in the post. For example, a relevant expert for a mobile phone advertisement post would be a 'tech specialist' and an irrelevant expert in this circumstance could be a doctor, and an influencer would just be labelled 'influencer'. The definition of expertise being used to determine the occupations used was 'a person with a high level of knowledge or skill relating to a particular subject or activity' (Cambridge Dictionary, no date). For all posts other than the vitamin advertisements, the irrelevant expert was represented by a doctor. This was because doctor was used as the relevant expert for the vitamin product, and using the same irrelevant occupation provided a means

of control to avoid bias toward some occupations over others. Expertise was presented to the right of the profile picture and above the product photograph either by itself or below the popularity indicator (see figure 5).

Product	Relevant expert	Irrelevant expert
Mobile phone	Tech specialist	Doctor
Water bottle	Water bottle tester	Doctor
Burger	Chef Food critic	Doctor
Vitamins/ supplements	Nutritionist Doctor	Maths Phd Astrophysicist
Journal	Stationary specialist	Doctor
Ear buds	Music producer	Doctor

Figure 5

The popularity attribute

Popularity was presented as the number of followers. Popularity was categorised either as higher or lower using a random number either <1000 as the low condition or >1 million as the higher condition. These brackets were chosen so that the difference between the higher and lower conditions of popularity was stark. The exact number of followers varied across the post bios, but the variation within either higher or lower popularity was far smaller than the variation between these two states. Similar to that of the attractiveness variable the popularity variable does not aim to categorise people as high or low popularity/ following as perceptions of what is high or low is subjective, hence two brackets were decided upon to allow there to be a clear difference between the higher and lower following conditions without making assumptions about how a certain number of followers may be perceived.

Higher popularity	1-3 million followers
Lower popularity	<1000 followers

Figure 6

Online platform and participants

The study was accessed by participants via the platform Amazon Mechanical Turk. Amazon Mechanical Turk is a crowdsourcing website through which researchers and businesses can pay 'workers' to carry out small online tasks or surveys. This platform was used as it allowed a large sum of participants to be obtained very quickly and allowed responses to be vetted to ensure that responses of participants that failed to follow the instructions could be removed before the analysis. The use of Amazon Mechanical Turk also allowed the privacy notice and information sheets to be easily linked at the beginning of the survey to ensure participants saw these documents before taking part. Each respondent was financially compensated for their time in accordance with the British national minimum wage for adults over the age of 23. Conducting the study online prevented any complications relating to COVID-19 transmission meaning there was no risk involved for the participants.

The posts were presented as tick box options on a Google Forms document linked through Amazon Mechanical Turk. Google Forms was used as the survey platform as it allowed for images to be used as response options within the same form as the regular Likert questions. The questionnaire was prefaced with the following statement; 'You must be able to name a social media influencer to take part in this questionnaire'. This was to ensure that participants could answer the 'influencer perception' part of the

survey. Despite this, there were a number of respondents that did not name an influencer when asked to in the survey. These participants' responses to the influencer perception section of the questionnaire were deemed unusable, however, if they had answered the experiment and self-esteem sections correctly their responses were still used for those sections.

Before choosing to partake Amazon Mechanical Turk workers could see information about what the survey involved so that they could choose whether they wished to take part. The information was as follows: 'Questions about your details (age and gender etc), self-esteem, and perception of social media influencers and products. At the beginning of the survey participants were told that: 'we are conducting a survey/ experiment investigating social media marketing preferences, perception of influencers and self-esteem, in relation to demographics such as gender, age and ethnicity. Please refer to the following documents before completion.' All Amazon Mechanical Turk users were allowed to partake in this experiment, however because of the demographic of the users the vast majority of participants were Americans aged 20-35.

Questionnaire design

The self-worth questions were kept identical to those of Rosenberg's 'self-esteem scale' so that the results could be comparable to other research findings using the same scale. However, the influencer perception section was deemed unusable due to the high volume of incorrect responses to the question requesting that a participant name a social media influencer, and evidence of inattention. Both sets of questions were answered using a 5-point Likert scale ranging from strongly disagree to strongly agree with 'neither agree nor disagree' in the middle with a score of 3. The 'neither agree nor disagree' option was added to allow the participant to not be unsure and not make false statements that might affect the validity of the data. The Rosenberg scoring system hence could not be used due to this extra value. Instead, a score was given to each participant by assigning points for each answer:

Questions 1, 2, 4, 6, 7 :

Strongly agree = 4

Agree = 3

Neither agree nor disagree = 2

Disagree = 1

Strongly disagree = 0

Questions 3, 5, 8, 9, 10:

Strongly agree = 0

Agree = 1

Neither agree nor disagree = 2

Disagree = 3

Strongly disagree = 4

The scale ranges from 0 – 40

Figure 7: Self-esteem scoring system

The participants self-esteem score could then be tested alongside their responses to the experiment section to test for a correlation between self-esteem score and which posts were 'liked' in the experiment.

Responses that exhibited 'straight-lining' in the influencer perception and self-esteem sections that utilised a Likert scale were deemed unusable for the experiment section as this indicated inattention. Straight-lining is a term used to describe a run of consecutive answers that are the same when it is very unlikely that they would be legitimately representative of the participants opinions. For example, choosing 'strongly agree' for every answer despite some of the statements being contradictory.

Information sheet :

https://docs.google.com/document/d/INuqmgAQjBJpfi1UGL-kPDjWX_Nkwv_JsHDFfxrEnuN8/edit?usp=sharing

Privacy Notice

: https://docs.google.com/document/d/17dpCwkILW0NdXF_tyKB8_2k_JcZFjQKa-7hqHpPlieo/edit?usp=sharing

At the end of the survey there was a debrief statement to avoid misleading participants:

'The social media posts presented in this study were created solely for the purpose of this research. The people in the images are not known to be social media influencers. The study aimed to investigate the influence of influencer attributes on cultural transmission within social media marketing'.

Analysis

Hypothesis I. Participants are more likely to click on common attributes than rare attributes.

Null: There is no significant difference between the frequency of likes for common attributes versus rare attributes.

In order to address this question a generalised linear mixed model was run in which 'generation' was treated as a numeric fixed effect and 'participant identification' as a random effect, with the likes for a specific attribute (1 representing a like for the higher variation and 0 representing a like for the lower variation) as the dependant variable. This produced a slope representing generation which could be exponentiated into an odds ratio. To determine significance of results $\alpha = 0.05$ will be used.

Popularity with attractiveness condition

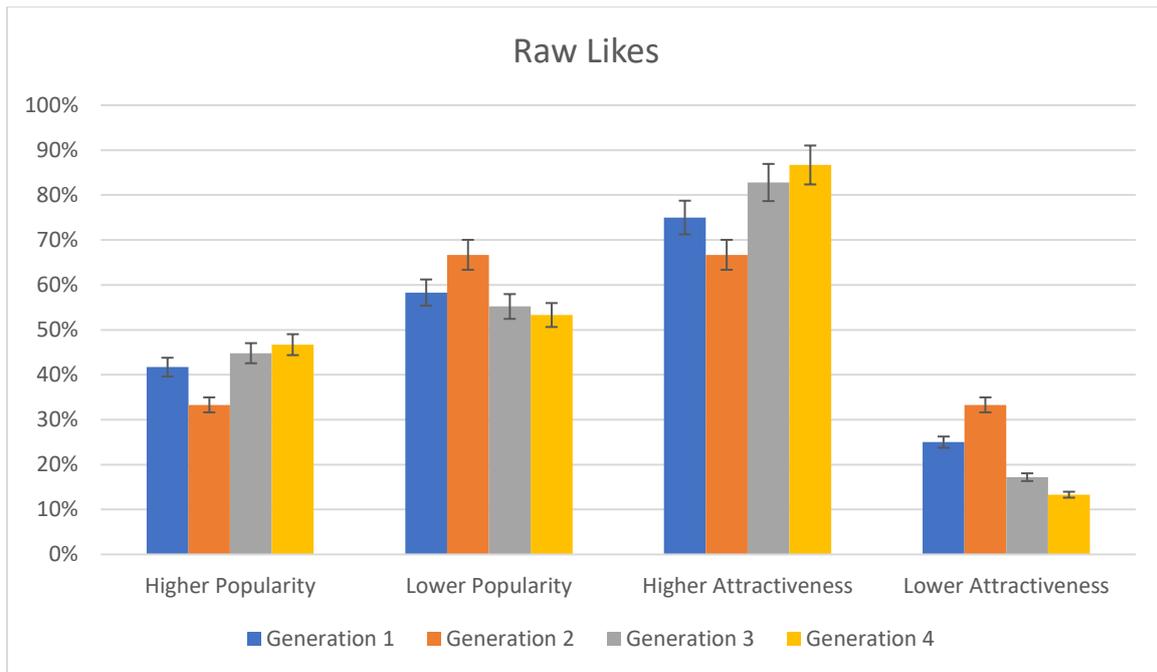


Figure 8: Bar chart showing like distribution in the popularity with attractiveness condition

Generation 1

	Higher popularity	Higher attractiveness
Proportion representation	50%	50%
Proportion of likes	41.7%	75%
Likes - representation	-8.3%	25%

The data suggest that there is an intrinsic preference in favour of posts displaying higher attractiveness rather than lower attractiveness, and a slight preference against higher popularity posts.

Generation 4

	Higher popularity	Higher attractiveness
Proportion representation	46%	79%
Proportion of likes	46.7%	86.7%
Likes - representation	0.7%	7.7%

The data suggest that there is a preference in favour of posts displaying higher attractiveness rather than lower attractiveness, and no significant preference towards higher or lower popularity posts in the final generation.

Popularity attribute

Odds of clicking a 'higher popularity' post = **0.5627049**

Probability of clicking 'higher popularity' post = 0.36008392

Attractiveness attribute

Odds of clicking a 'higher attractiveness' post = **5.196576**

Probability of clicking 'higher attractiveness' post = 0.83862055

Generalised linear mixed model with participant ID as a random effect and generation as a fixed effect;
Popularity attribute slope= 0.044

Odds ratio= **1.044982**

Attractiveness attribute slope= 0.469

Odds ratio= **1.598395**

The popularity attribute odds ratio suggests that there is no significant difference in the number of likes for higher popularity posts across generations, where $\alpha=0.05$. However, the odds ratio calculated for the attractiveness attribute implies that the greater the generation the higher the likelihood of a participant liking a higher attractiveness post. Hence, the frequency of likes increased to a level higher than would be expected from the intrinsic preference exhibited in the first generation. In the final generation 86.7% percent of participants' likes were for posts containing higher attractiveness, in comparison to the first generation where higher attractiveness posts were liked 75% of the time. This might suggest a possible effect of the transmission chain mechanism as likes increased with representation of the attractiveness variable however there was no significant difference in the likes for popularity across generations.

Attractiveness with expertise condition

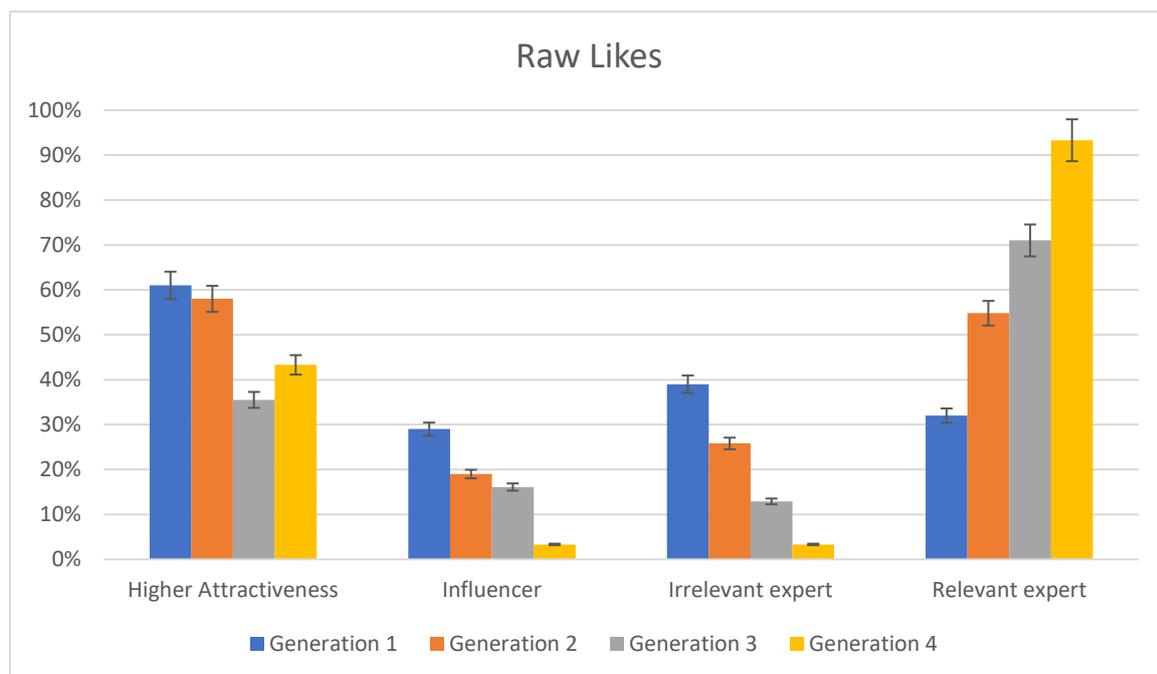


Figure 9: Bar chart showing like distribution in the expertise with attractiveness condition

Generation I

	Higher attractiveness	Irrelevant expert	Relevant expert	Influencer
Proportion representation	50%	33%	33%	33%
Proportion of likes	58.1%	19.4%	25.8%	54.8%
Likes - representation	8.1%	-13.6%	-7.2%	21.8%

The most significant preference in this condition is the preference for influencers, followed by the preference against irrelevant experts. There was also a small preference for posts exhibiting higher attractiveness and a small preference against posts featuring relevant experts.

Generation 4

	Higher attractiveness	Irrelevant expert	Relevant expert	Influencer
Proportion representation	42%	13%	71%	17%
Proportion of likes	23.3%	20%	73.3%	6.67%
Likes - representation	-18.7%	7%	2.3%	-10.33

The data suggest there is an intrinsic preference against higher attractiveness posts and influencers, as well as a slight preference for irrelevant experts. However, in the final generation influencers and higher attractiveness posts are less popular than in the first generation.

Generalised linear mixed model with participant ID as a random effect and generation as a fixed effect;

Attractiveness attribute slope= -2.396

Odds ratio= 0.09108155

Irrelevant expert slope= 0.078

Odds ratio= 1.081123

Relevant expert slope= 3.390

Odds ratio= 29

Influencer slope= -3.825

Odds ratio= 0.02181844

The odds ratios suggest a strong increase in the number of likes across generations for posts containing relevant experts and no significant effect of generation on likes for posts containing irrelevant experts. It can also be observed that there is a negative relationship between the number of likes for influencers and higher attractiveness across generations. This suggests that the higher representation of influencers following the first generation, did not influence participants in being more likely to choose a post containing an influencer despite the fact they were represented at least twice as much as relevant experts and irrelevant experts. Moreover, the odds ratio for relevant experts is high suggesting that increasing the generation increased the chances of a participant choosing a post containing a relevant expert, despite the fact that relevant experts had a low representation in the first and second generations.

Popularity only condition

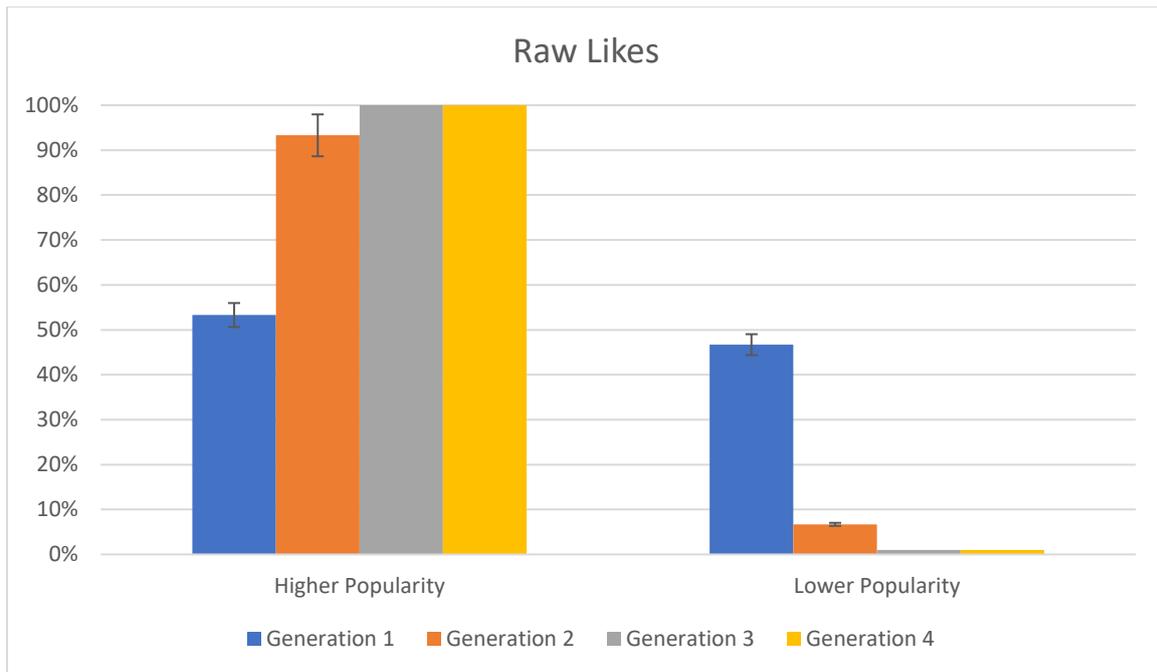


Figure 10: Bar chart showing like distribution in the popularity only condition

Generation 1

	Higher popularity
Proportion representation	50%
Proportion of likes	53.3%
Likes - representation	3.3%

Any preference for higher popularity or lower popularity was insignificant in this generation.

Generation 4

	Higher popularity
Proportion representation	100%
Proportion of likes	100%
Likes - representation	0%

Generalised linear mixed model with participant ID as a random effect and generation as a fixed effect;
 Popularity attribute slope= 4.200
 Odds ratio= 66.68633

The odds ratio produced in this condition is very high because one hundred percent of the fourth generation posts contained high popularity hence one hundred percent of participant responses were for posts containing high popularity. This would suggest that generation had a large impact on participant like choices. Observing the first generation it can be seen that there was no significant preference for higher or lower popularity. However, 91.7% of participant likes were for posts containing higher popularity in generation three hence why all of the posts in the fourth generation exhibited higher popularity. It must be noted that when the order of posts shown to the participant was generated for the third generation, coincidentally all apart from one of the higher popularity posts were presented in the first twelve posts, and all of the lower popularity posts apart from one were in the latter twelve. This meant that the higher popularity posts were seen by the participants first, perhaps introducing a form of bias in this condition.

Popularity with expertise condition

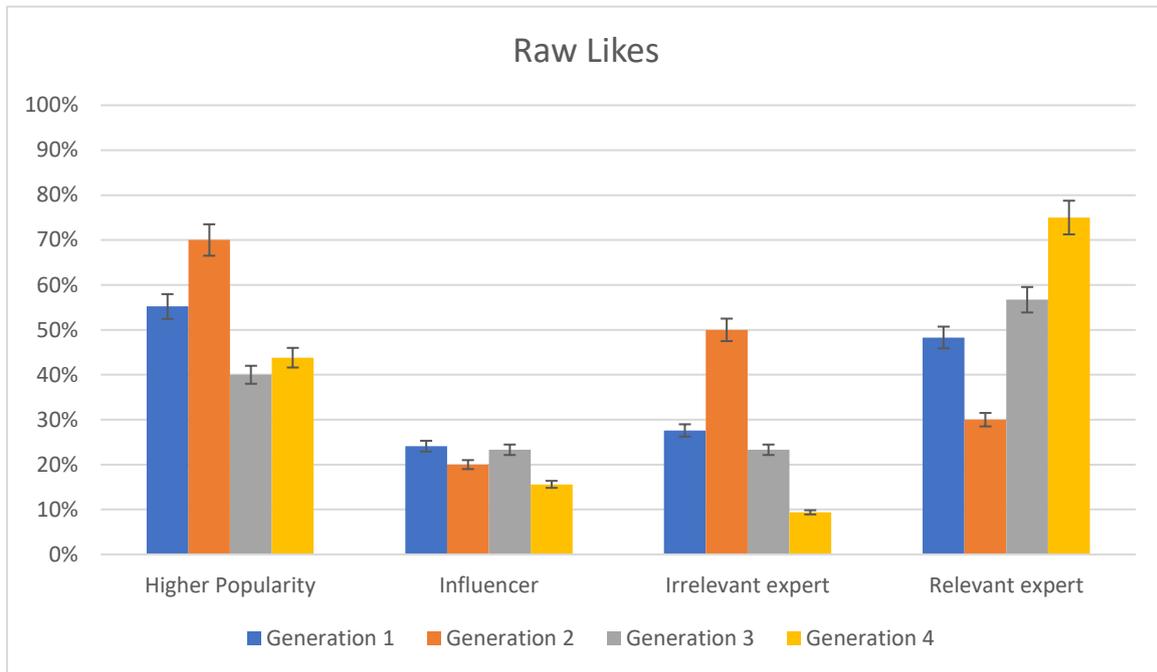


Figure II: Bar chart showing like distribution in the popularity with expertise condition

Generation I

	Higher popularity	Irrelevant expert	Relevant expert	Influencer
Proportion representation	50%	33%	33%	33%
Proportion of likes	55.2%	27.6%	48.3%	24.1%
Likes - representation	5.2%	-5.4%	15.3%	-8.9%

The most significant preference in this condition was the preference for relevant experts followed by the preference against influencers. There was also a very slight preference for higher popularity posts and a very slight preference against irrelevant experts.

Generation 4

	Higher popularity	Irrelevant expert	Relevant expert	Influencer
Proportion representation	42%	21%	54%	25%
Proportion of likes	43.8%	9.38%	75%	15.6%
Likes - representation	1.8%	-11.62	21%	-9.4%

Generalised linear mixed model with participant ID as a random effect and generation as a fixed effect;
 Popularity attribute slope= -1.107
 Odds ratio= 0.3305491
 Irrelevant expert= -1.685
 Odds ratio= 0.1854444
 Relevant expert= 2.586
 Odds ratio= 13.27656
 Influencer= -1.255
 Odds ratio= 0.2850758

The odds ratios suggest that there is a negative relationship between generation and the frequency of likes for higher popularity, irrelevant experts and influencers. However, there is a positive relationship between generation and the frequency of likes for relevant experts. As relevant experts were the most popular choice in the first generation, they were therefore represented at a higher frequency in the next generation meaning the high odds ratio may be evidence of the transmission chain mechanism affecting participants' choices as the frequency of likes in the final generation was higher than in the first.

Hypothesis 2. Higher popularity posts will have a greater like frequency than lower popularity posts and higher attractiveness posts will have a lower like frequency than lower attractiveness posts.

Null: There is no significant difference in the frequency of likes for the higher variations of popularity and attractiveness and the lower variations.

To test this a generalised linear mixed effects model was run in R, where the dependant variable is the attribute which can either be high or low (1 or 0), and participant identification has been accounted for as a random effect. Only the first generation of each condition was tested to avoid any possible bias that may result from the changing attribute representations. The probability of a participant choosing the higher variation of a certain attribute was calculated:

$$\text{Probability of clicking 'higher variation of attribute'} = \text{odds}/(1+\text{odds})$$

Popularity only condition (See figure 10)

Odds of clicking a post featuring higher popularity = 1.440514
 Probability of clicking a post featuring higher popularity = **0.67186542**

Attractiveness with popularity condition (See figure 8)

Popularity attribute
 Odds of clicking a 'higher popularity' post = 0.5627049
 Probability of clicking 'higher popularity' post = **0.36008392**

Attractiveness attribute

Odds of clicking a 'higher attractiveness' post = 5.196576
 Probability of clicking 'higher attractiveness' post = **0.83862055**

Attractiveness with expertise condition

Attractiveness attribute

Odds of clicking a 'higher attractiveness' post = 2.44489

Probability of clicking 'higher attractiveness' post = 0.70971497

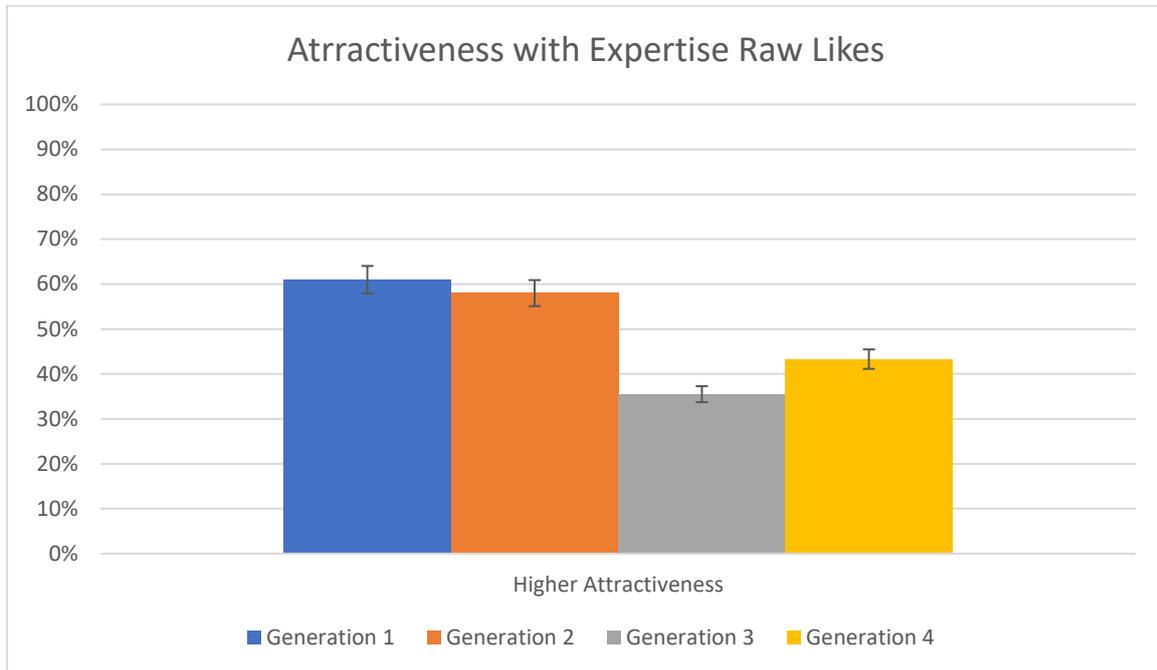


Figure 12: Bar chart showing attractiveness like distribution in the attractiveness with expertise condition

Popularity with expertise condition

Popularity attribute

Odds of clicking a post featuring higher popularity = 1.743684

Probability of clicking a post featuring higher popularity = 0.63552654

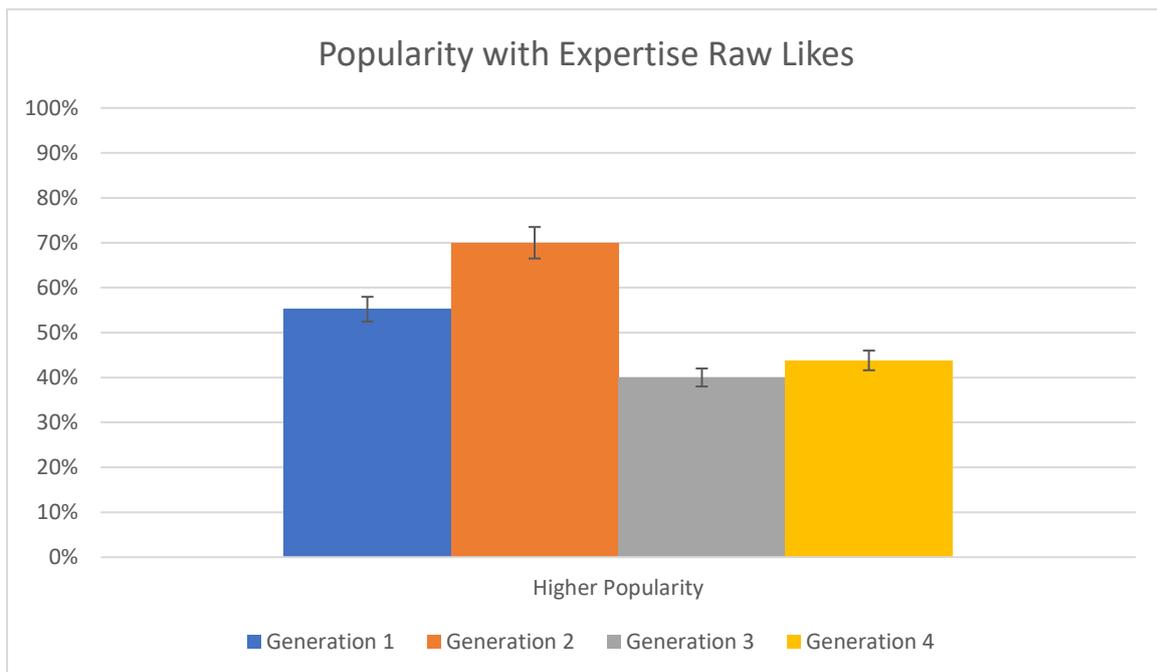


Figure 13: Bar chart showing popularity like distribution in the popularity with expertise condition

The probabilities calculated suggest that there is a preference for higher popularity over lower popularity as the probability of a participant liking a post exhibiting higher popularity in the first generation of the popularity only condition was around 0.67. Likewise higher popularity was preferred over lower popularity in the first generation of the popularity with expertise condition. Furthermore, it can be seen that there is a strong preference for higher attractiveness as in the first generation of the attractiveness with expertise condition the probability of a higher attractiveness post being chosen was around 0.71. However, in the attractiveness with popularity condition the probability of a person liking a higher popularity post was only 0.36 while the probability of liking a higher attractiveness post was approximately 0.84. Therefore, while there is a preference for higher attractiveness over lower attractiveness and higher popularity over lower popularity in three out of four conditions, lower popularity is favoured in the attractiveness with popularity condition.

Hypothesis 3. There will be no significant preference for neither higher attractiveness nor higher popularity when indicators of expertise are present.

Null: There will be no significant preference for neither higher or lower popularity nor higher or lower attractiveness when these attributes are presented in addition to indicators of expertise.

To test this a generalised linear mixed effects model was run in R, where the dependant variable is the attribute which can either be high or low (1 or 0), and participant identification has been accounted for as a random effect. Only the first generation of each condition was tested to avoid any possible bias that may result from the changing attribute representations. The probability of a participant choosing the higher variation of a certain attribute was calculated in the same way as for the previous hypothesis.

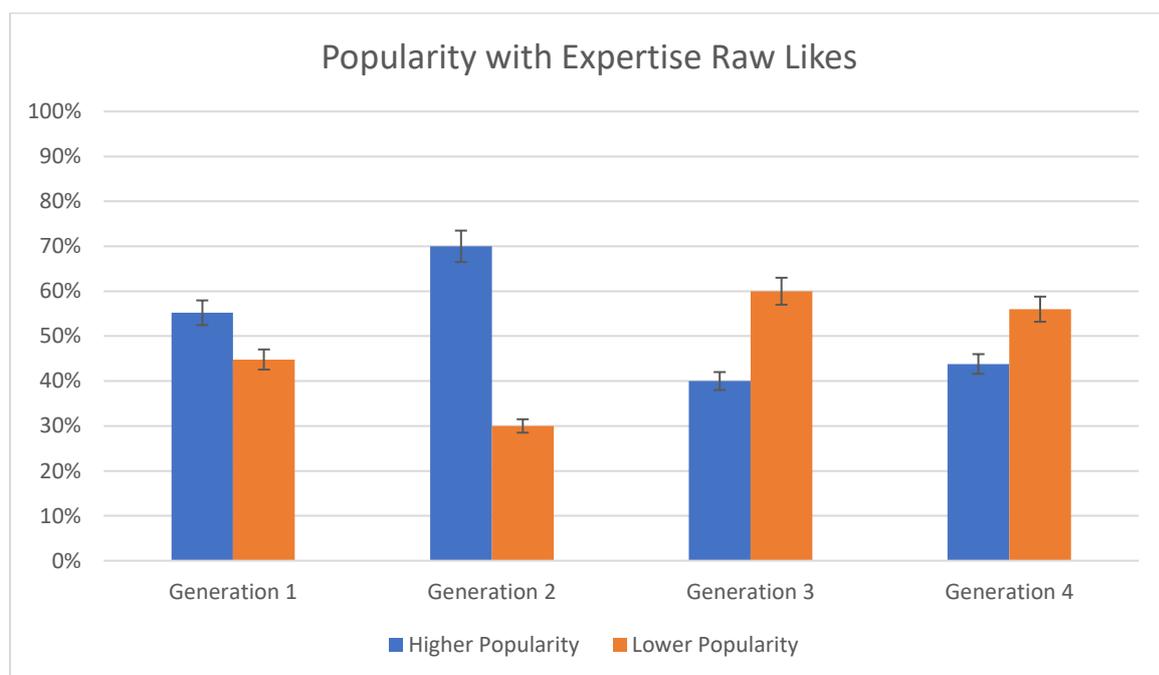


Figure 12: Bar chart showing popularity like distribution in the popularity with expertise condition

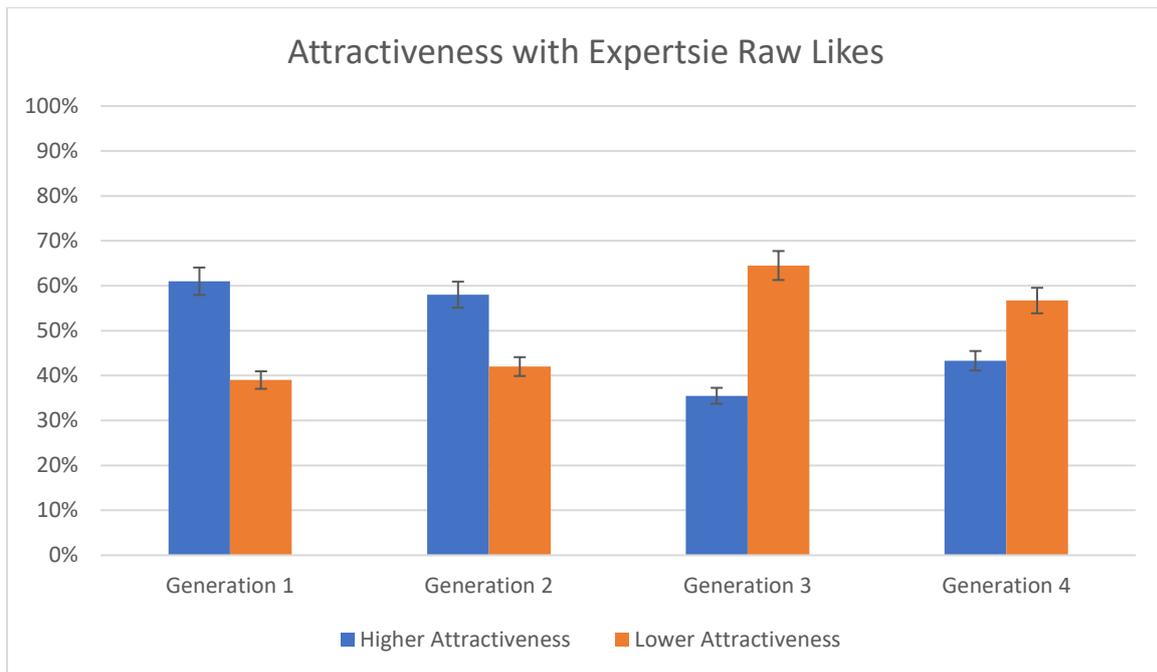


Figure 13: Bar chart showing attractiveness like distribution in the popularity with attractiveness condition

Popularity with expertise condition

Popularity attribute

Odds of clicking a post featuring higher popularity = 1.743684

Probability of clicking a post featuring higher popularity = **0.63552654**

Attractiveness with expertise condition

Attractiveness attribute

Odds of clicking a 'higher attractiveness' post = 2.44489

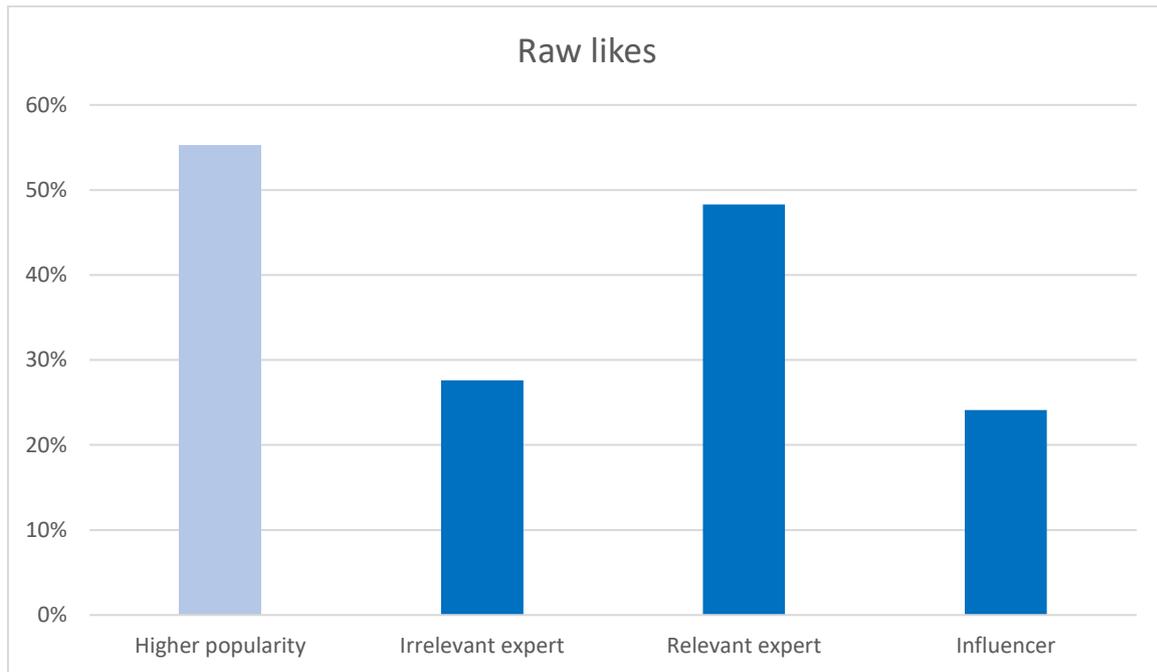
Probability of clicking 'higher attractiveness' post = **0.70971497**

It can be seen from the results of the tests for the second hypothesis that there was a preference for higher popularity posts in the first generation of the popularity only condition. When looking at the popularity with expertise condition in comparison to the popularity only condition, the likelihood of a participant liking a post containing higher popularity in both conditions was similar (0.037 difference). This suggests that the presence of expertise information did not significantly impact the participants' preferences for popularity. Looking at the attractiveness with expertise condition, it can be seen that there is a strong preference for attractiveness. However, when comparing the attractiveness with expertise condition, and attractiveness with popularity conditions, the probability of choosing a higher attractiveness post was 0.13 higher in the attractiveness with popularity condition than in the attractiveness with expertise condition. Hence, it is possible that attractiveness may be less important to participants when it is presented in addition to expertise information, however the preference is still significant, hence the null hypothesis can be rejected.

Hypothesis 4. Relevant experts receive more likes than irrelevant experts.

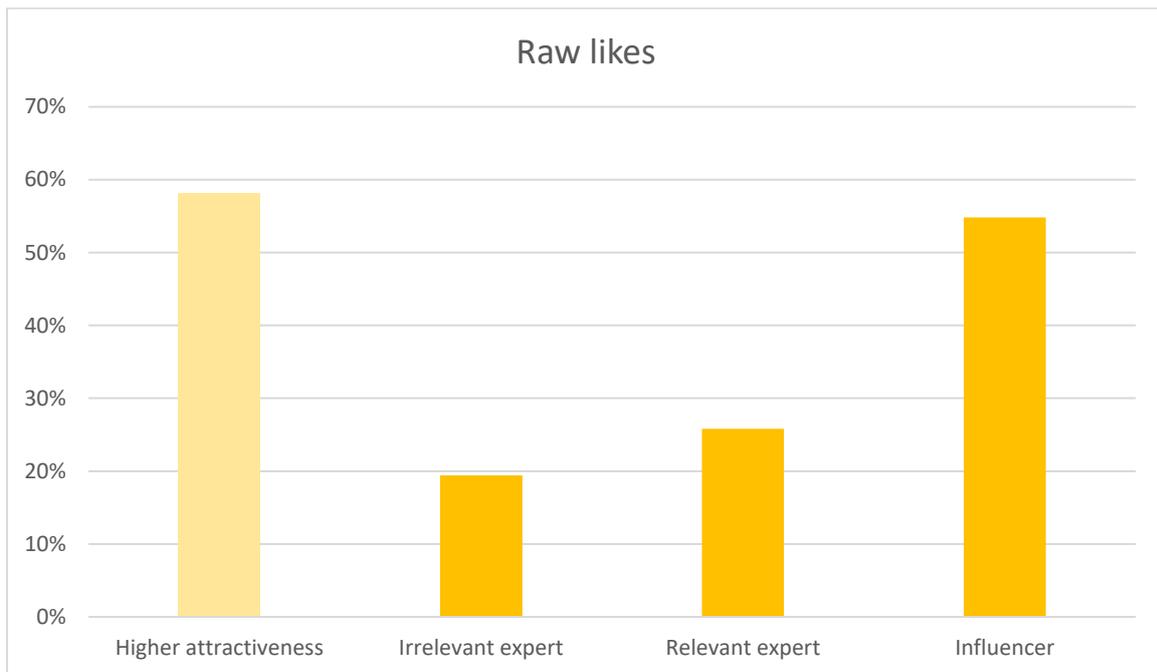
Null: There is no significant difference between the number of likes for relevant experts and the number of likes for irrelevant experts.

Popularity with expertise condition first generation



	Higher popularity	Irrelevant expert	Relevant expert	Influencer
Proportion representation	50%	33%	33%	33%
Proportion of likes	55.2%	27.6%	48.3%	24.1%
Likes - representation	5.2%	-5.4%	15.3%	-8.9%

Attractiveness with expertise condition first generation



	Higher attractiveness	Irrelevant expert	Relevant expert	Influencer
Proportion representation	50%	33%	33%	33%
Proportion of likes	58.1%	19.4%	25.8%	54.8%
Likes - representation	8.1%	-13.6%	-7.2%	21.8%

Results of a generalized linear mixed model

Attractiveness with expertise condition

Irrelevant expert

Intercept = -3.139

Odds of clicking a post featuring an irrelevant expert = 0.04

Probability of clicking a post featuring an irrelevant expert = **0.04**

Relevant expert

Intercept = -2.573
Odds of clicking a post featuring a relevant expert = 0.08
Probability of clicking a post featuring a relevant expert = 0.07

Popularity with expertise condition

Irrelevant expert
Intercept = -2.323
Odds of clicking a post featuring an irrelevant expert = 0.1
Probability of clicking a post featuring an irrelevant expert = 0.09

Relevant expert
Intercept = -0.186
Odds of clicking a post featuring a relevant expert = 0.83
Probability of clicking a post featuring a relevant expert = 0.45

In the attractiveness with expertise condition there was no significant difference between the number of likes for relevant experts and the number of likes for irrelevant experts. Whereas, in the popularity with expertise condition a participant was significantly more likely to like a post containing a relevant expert than a post containing an irrelevant expert.

Hypothesis 5. Those with a high self-esteem score are less likely to like the 'higher' variations of attributes than those with lower self-esteem scores.

Null: There is no significant relationship between participant self-esteem score and their choice of higher or lower variations of attributes.

Popularity with attractiveness condition

Attractiveness self-esteem odds ratio= 2.685857
Popularity self-esteem odds ratio= 1.440514

Popularity only condition

Popularity self-esteem odds ratio= 0.2424402

Popularity with expertise condition

Popularity self-esteem odds ratio= 0.1335873
Irrelevant expert self-esteem odds ratio= 0.1643101
Relevant expert self-esteem odds ratio= 6.951796
Influencer self-esteem odds ratio= 0.77958

Attractiveness with expertise condition

Attractiveness self-esteem odds ratio= 0.7032801

Irrelevant expert self-esteem odds ratio= 1.181754

Relevant expert self-esteem odds ratio= 0.8294437

Influencer self-esteem odds ratio= 1.246077

In the popularity with attractiveness condition the greater the self-esteem value of the participant the greater the odds of them liking a higher popularity or higher attractiveness post. In the popularity only condition the odds ratio suggests that a higher self-esteem score decreases the odds of a participant liking a higher popularity. In the popularity with expertise condition the effect of self-esteem score on attribute like preference depends on the attribute in question. For example, similar to the popularity only condition, in the popularity with expertise condition the higher the self-esteem score the smaller the odds of that participant choosing to like a post containing higher popularity. The odds ratios for influencer likes and irrelevant experts were also smaller than 1 in this condition further suggesting a decrease in likes for posts containing those attributes as self-esteem score increased. The only odds ratio greater than 1 in the popularity with expertise condition was that for the relevant experts, signifying that participants with higher self-esteem scores were more likely to choose relevant expert posts.

Finally, in the attractiveness with expertise condition the expertise odds ratios represented the opposite effect to those in the popularity with expertise condition, with ratios greater than 1 for irrelevant experts and influencers and less than 1 for relevant experts. Moreover, the attractiveness odds ratio was also less than 1 unlike in the attractiveness with popularity condition.

Results and discussion

Hypothesis I. Participants are more likely to click on common attributes than rare attributes.

The first hypothesis was designed to investigate whether attributes that were more common were more likely to be liked than attributes that were rare, in order to gauge if there was an effect resulting from the transmission chain mechanism. The transmission chain mechanism was inspired by the idea that social media recommends us content that is new to us but is already popular amongst others, hence this hypothesis aimed to test whether there is an inherent preference for attributes that are more common. However, it must be noted that participants were not aware that they were part of a transmission chain at the time of doing the experiment. In the attractiveness with popularity condition there was a preference for higher attractiveness posts which got stronger across generations, however there was no significant change across generations in likes for the popularity attribute. This positive relationship between attractiveness and generation supports the hypothesis that attributes that are common will receive more likes than attributes that are rare. Furthermore, in the popularity with expertise condition, relevant experts were liked the most in the first generation and there was a positive association between

generation and like frequency, meaning that the higher frequency of relevant expert posts may have encouraged more likes across the generations. Similarly, irrelevant experts and influencers received a lower like frequency in the first generation meaning they were represented less in the second generation, and the like frequency for these attributes decreased across the generations overall. The like frequencies for these attributes therefore expressed an overall preference for common attributes, and a preference against uncommon attributes. However, in the attractiveness with expertise condition influencers were liked the most in the first generation meaning they were represented over twice as frequently as relevant and irrelevant experts in the following generation. Despite the higher representation of influencers after the first generation the like frequency decreased across the following generations. Moreover, relevant experts had a low representation after the first generation however the like frequencies increased across the generations that followed. These observations therefore dispute the idea that common attributes receive more likes and uncommon attributes receive less.

Consequently, we cannot make an overarching statement about the effect of representation on like frequency, as the effect was different for the same attribute in different conditions. For example, higher representation led to higher like frequencies for the expertise attribute in the expertise with popularity condition, however this was not the case in the expertise with attractiveness condition. Therefore, the effect of generation on like frequencies is inconsistent, meaning further research must be done in order to determine the cause of the changing like frequencies over generations.

A further example of like frequencies increasing over generations was that of the popularity only condition. However, it was noticed upon reflection on the like count data that the higher variation posts were gathered in the first half of the group of posts presented to participants, and that across the different conditions the majority of likes were for posts presented in the first half of the group. For example, the most frequently liked post position was post number three, which varied in content across different conditions and generations, hence it is possible there was a subconscious preference for earlier appearing posts at play here. This could be an inherent bias toward the options we come across first, or it may be due to participant inattention or laziness meaning they did not look through all the posts before choosing which to like, as they were instructed to do. This could be an area in which to conduct further research as it may help to explain the changing like frequencies across generations that did not result from changing representation.

A way to mitigate the effects of the possible bias resulting from the order of posts could be to use a smaller number of posts so that the order becomes less significant as participants would be more likely to look through all of the options if there were fewer of them. Additionally, a like count could be added to each post to make clear which posts had been popular among previous participants and current participants would be made aware of the transmission chain mechanism. This would alter the dynamic of the experiment as the experiment would now be explicitly investigating the extent to which previous participants' responses affect current participants responses, rather than investigating whether there is an inherent preference for attributes that are common over attributes that are rare regardless of the subconscious reason as the participants were not aware of the transmission chain. However, it would be useful to compare the results from the modified experiment to the results of the original experiment to see if the responses changed. If the results showed a stronger relationship between representation and like frequency in the modified experiment than in the original then it might be that there is not an inherent preference for common attributes, rather a preference for posts that received greater deference and that this appreciation from others was not an assumed feature of the common posts in the original experiment.

Hypothesis 2. Higher popularity posts will have a greater like frequency than lower popularity posts and higher attractiveness posts will have a lower like frequency than lower attractiveness posts.

The hypotheses were devised in accordance with the concepts of cultural learning and prestige bias which suggest that humans are inclined to learn from individuals that they perceive to be prestigious in some way in order to learn adaptive information quicker. As a result, hypotheses numbers two and four were developed in order to test for preferences for prestigious attributes over non-prestigious attributes for example; higher attractiveness over lower attractiveness, higher popularity over lower popularity and relevant expertise over irrelevant expertise. Analysing the results from the second hypothesis testing the preferences for or against higher attractiveness and higher popularity, a strong preference for higher attractiveness can be seen, irrespective of the condition. This corroborates the literature on the attractiveness halo-effect for example, the Klebl et al. study concluding that attractive people are more highly associated with moral traits (2022) and Baker and Churchill who asserted that attractive people may be perceived to be more credible (2018). This finding would suggest that higher attractiveness may contribute to the success of influencer marketing strategies on social media, as consumers may be associating highly attractive individuals with credibility and appealing characteristics that encourage the consumer to think more highly of the product being advertised. Hence, it may be advised that companies choose attractive influencers to promote their products to increase engagement. However, the results from this experiment may not translate into reality as while the experiment was reminiscent of Instagram, participants knew it was only a simulation within a questionnaire therefore there were no 'purchase intentions' involved. Furthermore, the results from this study suggest that attractiveness may increase social media like counts but these do not necessarily correlate with purchase intentions.

Higher popularity was preferred in the popularity with expertise condition and the popularity only condition, however the probability of a participant choosing to like a higher popularity post in the popularity with attractiveness condition was only 0.36. This suggests an intrinsic bias towards higher popularity over lower popularity given the preference observed in the popularity only condition, however this preference appears to only stand in the absence of variation in attractiveness. It is possible that there is a stronger preference for attractiveness than popularity and that the presence of attractiveness indicators lessens the participants interest in popularity. This might suggest that social media product-endorsers may not need a significant following in order to gain engagement with their marketing posts providing they are attractive. However, further research would need to be done to fully understand this relationship of attractiveness relative to popularity.

Hypothesis 3. There will be no significant preference for neither higher attractiveness nor higher popularity when indicators of expertise are present.

The third hypothesis was designed to test the theory that humans prioritise expertise when choosing from whom to socially learn. Brand et al. asserted that prestige bias is only adopted in the absence of skill bias. Therefore, in the context of the experiment 'expertise' indicates a specific skill whether it be relevant to the product being advertised or not. Hence, if the statement by Brand et al. were correct we would have expected there to be no significant preference for popularity or attractiveness when experts were present. However, the results suggested that the presence of indicators of expertise did not negate the pre-existing preferences for the popularity or attractiveness attributes. While the probability of a participant choosing a higher attractiveness post was 0.13 lower in the attractiveness with expertise

condition than in the attractiveness with popularity condition the probability was still high hence the preference for attractiveness was still strong. Consequently, the null hypothesis suggesting that there would be no significant preference for neither higher popularity nor higher attractiveness in the expertise conditions, must be rejected. It is possible however, that the presence of expertise information mediated the preference for attractiveness and that participants may have viewed attractiveness as less important when expertise was present, though it was still important.

Hypothesis 4. Relevant experts receive more likes than irrelevant experts.

This hypothesis was designed to test for evidence of a generalised skill bias or for support of the theory that product-endorser fit increases the 'believability' of an endorsement, evoking a more favourable attitude toward the product (Kamins and Gupta, 1994, cited in Schouten et al., 2020). The results of the generalised linear mixed model suggested that there was no significant preference for relevant experts over irrelevant experts when attractiveness was present. However, there was a preference for relevant experts over irrelevant experts when an indication of popularity was present. Therefore, it is possible that while preferences for popularity, attractiveness and relevant experts were all observed in at least one condition, the strength of the preferences relative to one another may vary. For example, given the lack of preference for relevant experts in the attractiveness with expertise condition and the presence of this preference in the popularity with expertise condition, it is possible that the preference for attractiveness is stronger than the preference for expertise and the preference for popularity is not stronger than that for expertise. These results corroborate the theory resulting from the second hypothesis that participants seem to exhibit a preference for one prestigious trait only rather than multiple. Moreover, it is possible that attractiveness is the attribute that is prioritised as the results for the second hypothesis similarly suggested a preference for popularity only in the conditions in which attractiveness was controlled.

Hypothesis 5. Those with a high self-esteem score are less likely to like the 'higher' variations of attributes than those with lower self-esteem scores.

Barkow theorised that prestige seeking behaviours result from the desire to maintain self-esteem, therefore hypothesis 5 tested whether a participant's self-esteem score correlated with their behaviour toward prestigious traits (Barkow *et al.*, 1975, p.556). Furthermore, Bearden et al. asserted that consumers with low self-esteem are more likely to be influenced by others and might look for assistance when making purchasing decisions (Bearden *et al.*, 2001). If the results were to support these statements from the literature, we would expect lower self-esteem scores to correlate with higher chances of liking posts containing prestigious traits such as higher popularity, higher attractiveness, or expertise, as these people may in theory be less influenced by others as they have more confidence in themselves. The results in the attractiveness with popularity condition did not support the theory that those with higher self-esteem would be less likely to like posts containing prestigious traits, rather higher self-esteem scores correlated with greater chances of liking higher popularity and higher attractiveness posts. However, in the popularity only condition, higher self-esteem led to a decreased chance of liking a higher popularity post. Furthermore, the expertise conditions presented conflicting findings to one another making it difficult to draw any conclusions for this hypothesis as the results seem to depend completely on the condition and the attribute in question.

Conclusion

The transmission chain unique to the experiment within in this study enabled the responses from one participant to determine which posts the next was shown. However, there was no consistent effect of this process across conditions and attributes, as posts containing more common attributes were not consistently liked more or less frequently than posts containing less common attributes. The experiment could be altered to include fewer fake social media posts for participants to choose from in order to minimise bias that may have resulted from the order in which posts were presented. Moreover, the addition of a statement informing participants of the transmission chain process prior to answering the experimental section, may change the way that participants respond and hence change the impact of the transmission chain mechanism. It would be interesting to compare the results of the experiment with these alterations to the original results to try to uncover why like frequencies changed across generations and whether this was related to a subconscious bias toward content that is more common on social media.

This study has found a strong preference for attractiveness in the context of a social media simulation experiment. Participants chose to like higher attractiveness social media posts consistently more often than lower attractiveness posts irrespective of condition. This observation corroborates the literature surrounding the attractiveness halo-effect and suggests that more attractive social media product-endorsers may be more successful in fostering engagement than those that are less attractive. Popularity was also investigated in the form of a social media following, and it was found that there was also a preference for posts by individuals with a larger number of followers versus those with much fewer. However, this was not the case in the popularity with attractiveness experimental condition in which a greater frequency of likes for lower popularity endorsers was observed. This highlights a potential avenue for further research as it is suggestive of an attribute hierarchy whereby some prestigious traits are preferred only in the absence of other prestigious traits.

The experiment also showed that there was a preference for experts which were relevant to the product being advertised, over experts in an unrelated field when indicators of popularity were also present. However, there was no significant preference for relevant experts over irrelevant experts when variation in attractiveness was displayed as well. This supports the theory that attractiveness may be a trait that is prioritised by the participants. Furthermore, it was found that the preference for expertise did not negate the preferences for popularity and attractiveness seen in the conditions without indicators of expertise. Hence, disproving the hypothesis that expertise may be more important to participants than attractiveness or popularity when choosing which posts to like in the experiment. Finally, the impact of participant self-esteem was inconsistent with a positive relationship between self-esteem score and frequency of likes for prestigious traits in some conditions and the opposite in others. More research must be done in order to better understand the effects of participant self-esteem as well as attribute representation, as both these variables produced both positive and negative relationships with like frequencies that varied depending on the attribute and condition in question.

References

- Baker, M. J. and Gilbert A. Churchill, J. (2018) 'The Impact of Physically Attractive Models on Advertising Evaluations', <https://doi.org/10.1177/002224377701400411>, 14(4), pp. 538–555. doi: 10.1177/002224377701400411.
- Barkow, J. H. *et al.* (1975) 'Prestige and Culture: A Biosocial Interpretation [and Comments and Replies]', *Current Anthropology*, 16(4), pp. 553–572. doi: 10.1086/201619.
- Barkow, J. H., O'Gorman, R. and Rendell, L. (2012) 'Are the New Mass Media Subverting Cultural Transmission?', <https://doi.org/10.1037/a0027907>, 16(2), pp. 121–133. doi: 10.1037/A0027907.
- Bearden, W. O., Hardesty, D. M. and Rose, R. L. (2001) 'Consumer Self-Confidence: Refinements in Conceptualization and Measurement', *Journal of Consumer Research*, 28(1), pp. 121–134. doi: 10.1086/321951.
- Bearden, W. O., Netemeyer, R. G. and Teel, J. E. (1989) 'Measurement of Consumer Susceptibility to Interpersonal Influence', *Journal of Consumer Research*, 15(4), pp. 473–481. doi: 10.1086/209186.
- Brand, C. O. and Mesoudi, A. (2019) 'Prestige and dominance-based hierarchies exist in naturally occurring human groups, but are unrelated to task-specific knowledge', *Royal Society Open Science*, 6(5), p. 181621. doi: 10.1098/rsos.181621.
- Das, R. (2021) *Instagram Algorithm 2022: How To Conquer It*. Available at: <https://statusbrew.com/insights/instagram-algorithm/#instagram-explore-page-algorithm> (Accessed: 13 April 2022).
- EXPERTISE* / *English meaning - Cambridge Dictionary* (no date). Available at: <https://dictionary.cambridge.org/dictionary/english/expertise> (Accessed: 30 April 2023).
- Kamins, M. A. and Gupta, K. (1994) 'Congruence between spokesperson and product type: A matchup

hypothesis perspective', *Psychology & Marketing*, 11(6), pp. 569–586. doi: 10.1002/MAR.4220110605.

Klebl, C. *et al.* (2022a) 'Beauty Goes Down to the Core: Attractiveness Biases Moral Character Attributions', *Journal of Nonverbal Behavior*, 46(1), pp. 83–97. doi: 10.1007/S10919-021-00388-W/METRICS.

Klebl, C. *et al.* (2022b) 'Beauty Goes Down to the Core: Attractiveness Biases Moral Character Attributions', *Journal of Nonverbal Behavior*, 46(1), pp. 83–97. doi: 10.1007/S10919-021-00388-W.

Lee, J. A. *et al.* (2022) 'The Psychological Consequences of Envy on Instagram', *https://home.liebertpub.com/cyber*, 25(11), pp. 703–708. doi: 10.1089/CYBER.2022.0001.

Ma, D. S., Correll, J. and Wittenbrink, B. (2015) 'The Chicago face database: A free stimulus set of faces and norming data', *Behavior Research Methods*, 47(4), pp. 1122–1135. doi: 10.3758/S13428-014-0532-5.

Miller, A. G. (1970) 'Social Perception of Internal-External Control', *http://dx.doi.org/10.2466/pms.1970.30.1.103*, 30(1), pp. 103–109. doi: 10.2466/PMS.1970.30.1.103.

Neugaard, B. (no date) *halo effect / psychology / Britannica*. Available at: <https://www.britannica.com/science/halo-effect> (Accessed: 2 January 2023).

RUBIN, A. M., PERSE, E. M. and POWELL, R. A. (1985) 'LONELINESS, PARASOCIAL INTERACTION, AND LOCAL TELEVISION NEWS VIEWING', *Human Communication Research*, 12(2), pp. 155–180. doi: 10.1111/J.1468-2958.1985.TB00071.X/ABSTRACT.

Santora, J. (2021) *Influencer Marketing Stats | 100 Influencer Marketing Statistics for 2021*. Available at: <https://influencermarketinghub.com/influencer-marketing-statistics/> (Accessed: 11 January 2022).