THIS MANUSCRIPT HAS BEEN ACCEPTED FOR PUBLICATION AT BEHAVIORAL
RESEARCH METHODS

**Do you know the *Woolly Bully*? Testing Era-Based Knowledge to Verify Participant Age
Online**

Rachel Hartman^{1,2*}, Aaron J. Moss^{1*}, Israel Rabinowitz¹, Nathaniel Bahn¹, Cheskie
Rosenzweig^{1,3}, Jonathan Robinson^{1,4}, Leib Litman^{1,5}

¹CloudResearch

²Department of Psychology and Neuroscience, University of North Carolina at Chapel Hill

³Department of Clinical Psychology, Columbia University

⁴Department of Computer Science, Lander College

⁵Department of Psychology, Lander College

*denotes equal contribution

6/23/2022

Wordcount: 7,326

Author Note

Correspondence concerning this article should be addressed to Leib Litman, Department of Psychology, Lander College, 75-31 150TH ST, Kew Gardens Hills, NY, 11367. Email: leib.litman@touro.edu

Abstract

People in online studies sometimes misrepresent themselves. Regardless of their motive for doing so, participant misrepresentation threatens the validity of research. Here, we propose and evaluate a way to verify the age of online respondents: a test of era-based knowledge. Across six studies ($N = 1,543$), participants of various ages completed an age verification instrument. The instrument assessed familiarity with cultural phenomena (e.g., songs and tv shows) from decades past and present. We consistently found our instrument discriminated between people of different ages. In Studies 1a and 1b, self-reported age strongly correlated with performance on the instrument (mean $r = .8$). In Study 2, the instrument reliably detected imposters who we knew were misrepresenting their age. For impostors, self-reported age did not correlate with performance on the instrument ($r = .077$). Finally in Studies 3a, 3b, and 3c, the instrument remained robust with African Americans, people from low educational backgrounds, and recent immigrants to the U.S. Thus, our instrument shows promise for verifying the age of online respondents, and, as we discuss, our approach of assessing “insider knowledge” holds great promise for verifying other identities within online studies.

Keywords: Amazon Mechanical Turk, data quality, age verification, online research

Do you know the *Wooly Bully*? Testing Era-Based Knowledge to Verify Participant Age Online

Wooly Bully is a quirky song originally recorded by Sam the Sham and the Pharaohs. In the song's lyrics a conversation takes place in which Mattie tells Hattie about a thing she saw—a thing with two big horns and a wooly jaw. If you were alive when the song came out in 1965, you probably heard *Wooly Bully* on the radio and possibly even danced to it. If, however, you were born in the decades following the 1960's, you are increasingly unlikely to know the song. You may have heard *Wooly Bully* in a movie or on the "oldies" station, but you are unlikely to know who sang the song, that the record sold three million copies worldwide, or that it was the number one song of 1965 despite never occupying the top spot on the Billboard charts. To know these details, it helps to have lived them.

The idea that people are more familiar with things that are part of their lived experience than things that are not is so obvious it is often taken for granted. Nearly everyone, for example, would expect a master brewer to know more about hops, grains, yeast, and the process of brewing beer than a guy at the bar drinking beer. Similarly, most people would predict that a person from Africa knows more about African geography than a person from Australia. And, at a group level, most people expect millennials to be more tech-savvy than baby boomers because millennials grew up with recent technology whereas baby boomers did not. While the idea of group differences in knowledge based on experience may be relatively mundane, the application of this idea may hold value for behavioral scientists who conduct online research. Specifically, group differences in the relative knowledge people possess may be a way to detect when participants misrepresent themselves in online studies (e.g., Kramer et al., 2014).

Participant misrepresentation in online research can take various forms, but sometimes manifests as fraud. Fraud occurs when people lie about their demographic characteristics to qualify for a study that they would otherwise be ineligible to complete. The motivation for this behavior is presumably to collect the study's compensation (e.g., Chandler & Paolacci, 2017; Wessling et al., 2017). In past research, participants have been found to misrepresent their sexual orientation, lie about their gender and age, claim to own pets or consumer products they do not actually own, claim experience with fictional items that do not actually exist, and to generally say or do anything necessary to gain access to studies that leave the door open for fraud (e.g., Chandler & Paolacci, 2017; Gadiraju et al., 2015; Kan & Drummey, 2018; MacInnis et al., 2020; Siegel et al., 2015; Wessling et al., 2017).

Although it can seem like everyone online is willing to lie, the overall percentage of people who misrepresent themselves in online participant platforms to gain access to studies may actually be quite small—a number in the single digits (Chandler & Paolacci, 2017; MacInnis et al., 2020). However, even a small percentage of liars within an entire platform can result in study samples that are filled with fraudulent respondents. This is because studies that target hard-to-reach groups typically take longer to complete and pay more money, presenting a target for fraud (Chandler & Paolacci, 2017; Wessling et al., 2017). If the target population exists only in small numbers and honest participants won't lie, then the only people left to take the study are imposters. As one example of this phenomena, consider Siegel and Navarro (2019) who found that participants on Amazon Mechanical Turk (MTurk) misrepresented themselves as Democrats or Republicans to gain access to studies targeted toward those groups. In the final sample of Democrats, 17% of participants were imposters whereas 45% of the Republican sample were

imposters. Because Republicans are a harder-to-reach group on MTurk than Democrats, the study targeting the harder to reach group wound up with more misrepresentation.

In the context of age misrepresentation, Wessling et al. (2017) found that many participants in a survey about dietary fiber supplements were dishonest. When participants were incentivized to do so, nearly half claimed to be over 50 years of age after providing a younger age in a previous survey. Because the average age of MTurk participants is about 35 (SD = 12; Burnham et al., 2018), people over age 50 are in limited supply. Imposters like this pose a serious threat to the validity of research not only because they are outside the target population, but also because if they go undetected, they are likely to systematically skew the findings of studies. For example, in Wessling et al.'s study (2017), young imposters expressed more interest in fiber supplements than adults truly over the age of 50 did (Chandler et al., 2020; Wessling et al., 2017). Similarly, men who impersonate women choose pink consumer products at a far higher rate than women themselves choose (Wessling et al., 2017). In these examples, the results were skewed toward overrepresentation of the interests of the target groups, but results can also be skewed in the opposite direction. For example, in Siegal and Navarro's (2019) study, Democrats posing as Republicans tended to underreport favorability towards Donald Trump.

To combat such fraud, there are several methodological steps researchers can take. These include separating demographic screeners from the actual study (e.g., Wessling et al., 2017), using a platform's demographic targeting options, and building a panel of participants for repeated survey use (e.g., Wessling et al., 2017). Yet, even when these best practices are employed, imposters may slip through because screening relies on participants' self-reported demographics (see MacInnis et al., 2020). An opportunistic millennial might, for example, realize that studies targeting baby boomers consistently offer higher pay. This person may then

create a user profile in which they consistently misrepresent their age. Moreover, outside of paid participant platforms, researchers sometimes recruit participants from open sources like Reddit or Craigslist (e.g., Antoun et al., 2016; Shatz, 2017), or intercept people while they play video games or shop online. When recruitment occurs in open sources, researchers lack the benefit of even basic demographic data that people typically provide to participant platforms and social media sites. Thus, whether there is a built-in mechanism for targeting people of specific ages or not, researchers may want a tool to further verify that participants who are recruited from hard-to-reach demographic groups—and who often are recruited at a premium—are who they say they are. It is for these instances that we propose examining relative group differences in knowledge to verify the characteristics of online respondents.

In this paper, we demonstrate a general approach to verifying participant characteristics in online studies using age as an example. In addition to the differences in lived experience we highlighted at the outset, there are good theoretical reasons to expect era-based differences in knowledge between younger and older adults. Research consistently finds that people form especially enduring memories for music, current events, and life experiences that occur in late adolescence and early adulthood (e.g., Beier & Ackerman, 2001; Janssen et al., 2008). Collectively, enhanced recall for this period of life is known as the “reminiscence bump.” While research on the reminiscence bump initially focused on autobiographical memories, more recent research indicates that it extends to semantic knowledge as well (see Janssen et al., 2008). Thus, given the strong association between the reminiscence bump and recall for music and events from one’s youth (Zimprich, 2020), we expected adults in their 50’s, 60’s, and 70’s to possess more knowledge about popular culture from past decades than people in their 20’s, 30’s, and 40’s. Conversely, we expected younger adults to possess more knowledge than older adults

concerning recent trends in pop culture, music, and television. This expectation was based on the well-established finding that recent events are remembered better than remote events, particularly among young adults (e.g., Kahana et al, 2002). Thus, we expected younger people to be both more interested in and aware of recent cultural phenomena, and to form better memories for those more recent events. Together, we expected the difference in people's knowledge of historical and contemporary culture to be especially diagnostic of their age and thus to aid in verifying age in online research.

Overview

We report the results of six studies that investigated whether people's relative knowledge of cultural phenomena can be used to determine their age. We begin by describing the development of materials for the Age Verification Instrument (AVI). In Studies 1a and 1b we report how well our instrument did in separating online respondents by self-reported age from two different participant recruitment platforms. In Study 2, we conducted a "stress test" of our instrument. After inviting younger adults to participate in a study that was advertised for people "50 years of age or older," we examined what percentage of respondents were willing to lie to take the study and how well our test of era-based knowledge did at categorizing the self-reported age of imposters. Finally, in Studies 3a, 3b, and 3c, we examined the generalizability of our test with participants of various demographic groups within the U.S. Across all studies, we predicted that younger people would know more about recent cultural phenomena than phenomena from decades past and that the opposite would be true for older adults. Because we expected *the difference* in era-based knowledge to be the main discriminator between younger and older adults (rather than absolute era-based knowledge), we used difference scores in our analyses. We have reported all measures, conditions, and data exclusions. Post-hoc power analyses revealed

that our sample sizes across studies provided 100% power to detect large effects. All data, study materials, and supplemental materials are available at:

https://osf.io/bn4xy/?view_only=7252e963f3bd4c0c981eed6ddd085ee8.

Instrument Development

Method

Participants and Design

Twenty adults from MTurk participated in the pilot study. We used CloudResearch (formerly TurkPrime; Litman et al., 2017) to sample people within the U.S. and to recruit people in different age groups. Specifically, we recruited ten people between the ages of 18 and 25 and ten people aged 65 or older. Data collection occurred in March 2019. People were paid \$2.50 for the survey which we expected to take about 25 minutes. Data collection ended after two hours.

Procedure

To determine which items best discriminate between older and younger adults, we created a list of multiple-choice questions about 69 popular songs and 20 TV shows from various decades. We provided participants with a list of names and asked them to identify the artist who recorded the songs or the main characters from the TV shows. There was also an option to select “I don’t know.” See the supplemental materials for more details about the curation process and the full list of items.

We further asked 14 open-ended questions about national politics and news stories (e.g., “Who is H.R. Haldeman or which major event was he connected to?” and “What is the name of the volcano that erupted in 1980?”). We also presented participants with headshots of 11 U.S. presidents and 11 first ladies and asked them to name the person pictured in an open-ended format. Finally, using open-ended questions, we asked participants when they graduated from

high school and how old they were during the Watergate scandal (participants under 50 typically wrote “not born yet”). We asked these open-ended questions for exploratory purposes and do not discuss them further here.

At the beginning of the survey, people read a message telling them not to use outside sources and that it was better to select “I don’t know” as an answer rather than to guess. The statement told people that compensation did not depend on performance.

Results

Analytic Approach

We sought to create a test of era-based knowledge that discriminates between younger and older adults by assessing the *relative knowledge* people possess. To determine the questions with the highest potential to discriminate between groups, we used the ‘difference method’. For each item in the pilot study, we subtracted the percent of people in the 18–25-year-old cohort who correctly answered the question from the number of people in the 65+ cohort who correctly answered the question, yielding a difference score. Positive difference scores indicated questions where older people were more likely than younger people to answer correctly and negative difference scores indicated questions where younger people were more likely than older people to answer correctly. Overall, we suspected that the difference score method would offer greater validity than separately analyzing items that represent past and present culture because these scores assess total era-based knowledge. By contrast, difference scores assess people’s relative knowledge. In other words, we examined whether people knew more about historical culture than contemporary culture, or vice versa, regardless of how much they actually know about culture.

Figure S1 in the supplemental materials shows the difference score for each question in the pilot study. The questions with the most positive and negative difference scores did the best job at discriminating between younger and older people. To select items for the final questionnaire, we chose 9 songs with the highest positive scores, 6 songs with the lowest negative scores, and 4 TV shows with the highest positive scores. See Table S1 for the 19-item AVI.

Study 1: Verifying Age Differences on MTurk and Prime Panels

In Study 1a, we tested the AVI's ability to discriminate between online participants of various ages using MTurk. We recruited people into the study by age based on CloudResearch's demographic targeting. We aimed to have approximately equal numbers of participants in each age group, and to thus assess the ability of our items to discriminate amongst online participants of various ages. In Study 1b, we replicated our results using Prime Panels, an aggregator of online panels commonly used in market research (see Chandler et al., 2019). Participants on Prime Panels are more representative of the U.S. population, especially in terms of age, providing us with better access to participants above age 50 (Chandler et al., 2019; Litman et al., 2020a).

Method

Participants and Design

Study 1a. Three hundred and two adults from MTurk participated in Study 1a. We used CloudResearch's MTurk Toolkit (Litman et al., 2017) to target participants within the U.S. and to recruit participants in different age groups. Specifically, we recruited 50 participants in six separate groups, with each group corresponding to a different decade of age (20s through 70s).

Participants were paid 50 cents to complete the study which we estimated would take three minutes. All data were collected in April 2019, and data collection ended after three days.

Study 1b. We recruited 350 adults from Prime Panels. As with Study 1a, we split the sample into six groups of approximately 50 participants each, with each group corresponding to a different decade of age. Because Prime Panels aggregates several panels to collect large samples, participants were compensated based on the platform they were recruited through. Some participants may have completed the study in exchange for flight miles, points, money, or other rewards. All data were gathered in April 2019; data collection closed after three hours.

Procedure

We presented participants with the AVI. Some of the questions had four response options and some had five. We instructed participants to answer to the best of their ability without using outside sources and to select “I don’t know” when applicable. We also stressed that we would not penalize participants if they did not know the answers. As in the instrument development study, we asked participants to provide information about their age to verify that the database information was accurate. These included open-ended questions about participants’ current age, the year they graduated from high school, and how old they were during Watergate (participants under 50 typically wrote “not born yet”).

For exploratory purposes we also asked participants to select the decade of their life in which they were the happiest and to elaborate on what was positive about that time. They then selected the decade of their life that was most difficult and described what made it so. The results from these items are not reported here.

Analytic approach

We used the difference method to assess each person's relative knowledge of historical (questions about pop culture prior to the year 2000) and contemporary (questions about pop culture after the year 2000) culture. For each participant, we separately summed the number of correct responses on all items measuring historical and contemporary knowledge, then converted the sums to percentages. Finally, we calculated a difference score by subtracting the percentage of correct responses to contemporary questions from the percentage of correct responses to the historical questions. This yielded a difference score variable with a range of -100% (correctly answered all contemporary questions and no historical questions) to +100% (correctly answered all historical questions and no contemporary questions).

For study 1a, in addition to using the CloudResearch database to target participants whom we expected to fall into six age groups, we asked participants to self-report their age. We opted to rely on self-reported age in our analyses because that is the data most researchers would have access to. There was a strong correspondence between self-reported age and database age ($r = .965, p < .001$).

In both studies we used linear regression, predicting the continuous self-reported age variable using performance on the AVI items as the predictor. We also assessed the utility of the instrument for distinguishing between decades of age. To do so, we split participants' self-reported age into six groups, with each group corresponding to a different decade, and tested differences between the groups using a one-way ANOVA. To assess the value of using the difference score, as opposed to just using the scores on the contemporary or historical questions, we also examined the correlations between self-reported age and each of these three measures.

Here and in later studies we tested both the full 19-item instrument (see Table S1) as well as a shorter 6-item subset (AVI-S). In addition to being easier to implement, research indicates

that much of the age-related differences in people's knowledge can often be captured by a few items rather than by multiple items (e.g., Schroeders et al., 2021). The 6-item subset comprised the three historical items that older adults most often answered correctly and the three contemporary items that younger adults most often answered correctly (see Table 1)¹. Historical items include Bonanza (1959-1973), The Way We Were (1974), and The First Time I Ever Saw Your Face (1969), while contemporary items include Somebody That I Used to Know (2011), How You Remind Me (2001), and Boom Boom Pow (2009). Most analyses showed the difference between the full 19-item scale and the shorter six-item version was insubstantial (e.g., the full scale predicted 68.2% of the variance in age, while the shorter version predicted 64.3%). Therefore, we report the results of the shorter scale. Analyses using the full version of the AVI are available in the supplemental materials.

Table 1

The Short Age Verification Instrument (AVI-S)

Item (decade)	Answer choices	Score
Bonanza (1959-1973)	a. Mark Lewis b. Shirley Dawson c. Ben Cartwright	Older cohort: 100% Younger cohort: 0% Difference score: 100%
The Way We Were (1974)	a. Barbra Streisand b. Freddie Hart c. Waylon Jennings d. Addrisi Brothers	Older cohort: 100% Younger cohort: 0% Difference score: 100%
The First Time I Ever Saw Your Face (1969)	a. The Elton John Band b. Roberta Flack c. Spiral Staircase d. The Peppermint Rainbow	Older cohort: 100% Younger cohort: 10% Difference score: 90%
Somebody That I Used to Know (2011)	a. Thompson Square b. Gotye c. Cher Lloyd	Older cohort: 30% Younger cohort: 70% Difference score: -40%

¹ We thank Reviewer 1 for pointing out the utility in analyzing a shorter scale

	d.Pitbull Featuring John Ryan	
How You Remind Me (2001)	a.Lifehouse b. Nickelback c.AFI d.Hinder	Older cohort: 0% Younger cohort: 40% Difference score: -40%
Boom Boom Pow (2009)	a.Lil Wayne Featuring Drake & Future b.Orianthi c. The Black Eyed Peas d.The Script	Older cohort: 10% Younger cohort: 70% Difference score: -60%

Note. Correct answers are bolded. All answer choices also included “I don’t know” as an option. Scores represent the percent of participants in the pilot study who answered the question correctly.

Results

Instrument Performance by Age

We ran a linear regression analysis predicting self-reported age with the AVI-S. The AVI-S significantly predicted self-reported age (Study 1a: $r = .80, p < .001$; Study 1b: $r = .75, p < .001$) and explained 64.3% (1a) and 58.5% (1b) of the variance in age.

Next, we tested differences between the group means using a one-way ANOVA with age groups as the predictor and the difference score as the dependent variable. The ANOVA was significant in both Study 1a ($F(5, 296) = 121.34, p < .001$), and Study 1b ($F(5, 362) = 113.6, p < .001$), see Figure 1. We followed up with post-hoc Tukey tests to determine which means differed.

In both studies, the scores of participants in the two youngest age groups were not significantly different from each other. Participants in their 40’s differed significantly both from younger participants and from older participants. Participants in their 50s scored higher than those in younger age groups but did not differ from participants in their 60s and 70s. See Table S2 (Study 1a) and Table S3 (Study 1b) in the Supplemental Materials for group descriptives.

Figure 1

Study 1 AVI-S Difference Scores by Age Group

Note. AVI-S = Age Verification Instrument (Short). Error bars represent 95% confidence intervals.

We next examined whether self-reported age was more closely related to the difference score or to the contemporary or historical items. For Study 1a and Study 1b, respectively, the continuous self-reported age variable had weaker correlations with contemporary items ($r = -.62$, $p < .001$) and ($r = -.55$, $p < .001$), and with historical items ($r = .74$, $p < .001$) and ($r = .66$, $p < .001$) than the difference score ($r = .80$, $p < .001$) and ($r = .75$, $p < .001$).

We additionally wanted to see how accurate the AVI-S was at categorizing people into their age groups. Using a simple categorization scheme, we grouped people into “older” and “younger” groups based on whether their score on the instrument was higher than zero or less than or equal to zero. In Study 1a, of the participants in their 50s, 60s, and 70s or over, 87.5%, 94.4%, and 95.5%, respectively, were correctly categorized as “older”. Participants in their 40s were evenly split between the “older” (56.9%) and “younger” (43.1%) groups. Of the participants in their 30s and 20s, 94% and 94.2% were correctly categorized as “younger”.

In Study 1b, of the participants in their 50s, 60s, and 70s or over, 82.5%, 97.6%, and 95.6%, respectively, were correctly categorized as “older”. Participants in their 40s were evenly split between the “older” (50.9%) and “younger” (49.1%) groups. Of the participants in their 30s and 20s, 78.4% and 96.4%, respectively, were correctly categorized as “younger”. Thus, on average, the AVI-S is about 92% accurate at categorizing people as being older than 50 or younger than 40. Table 2 presents a summary of these results.

Table 2

Precision Measures Across Studies

Study	Sensitivity	Specificity	Accuracy	Positive Predictive Value
Study 1a	91.3	93.1	92.0	95.1
Study 1b	92.5	86.0	89.8	90.5
Study 2*	NA	81.9	81.9	NA
Study 3a	88.6	87.9	88.3	89.4
Study 3b	89.0	89.7	89.4	89.9
Study 3c	74.3	82.4	78.3	80.6

Note: The measures in this table reflect the ability of the AVI to categorize participants as either “younger” (below age 41) or “older” (above age 50). *Sensitivity* refers to the ability of the AVI to correctly identify participants as older. *Specificity* refers to the ability of the AVI to correctly identify participants who are younger. *Accuracy* refers to the AVI both correctly classifying younger participants as younger as well as older adults as older. *Positive Predictive Value* refers to the ratio of participants correctly classified by the AVI as older relative to all participants classified by the AVI as older. All numbers represent percentages out of the sample for each study.

*For Study 2 with imposters, we used the age of participants stored in the CloudResearch database. All participants were selected because their age in the database was below 30, thus placing them into the “younger” category. *Sensitivity* and *Positive Predictive Value* are not relevant here because there were no older adults in the sample.

Discussion

In these initial studies, our goal was to verify that the instrument we developed was able to differentiate between younger and older adults. Using both MTurk and Prime Panels samples, we obtained strong support for our instrument. Participants 40 and under knew relatively little about historical culture and were relatively more knowledgeable about contemporary culture. Conversely, participants over 50 knew little about contemporary culture, but demonstrated greater knowledge about historical culture. Participants in their 40s fell in between, reflecting their shared knowledge of both historical and contemporary culture. Further, the instrument was highly effective at categorizing people into age groups: across both samples, roughly 92% of participants could be accurately categorized into age groups based on their score.

In Study 1 we validated the AVI-S with samples that had no incentive to lie about their age, since we recruited them for studies targeted to their age group. In Study 2 we provide a “stress test” of our questionnaire, testing to see if it can detect participants who lie about their age.

Study 2: Detecting Imposters

In Study 2, we wanted to examine the ability of our instrument to detect participants who we knew were lying about their age, i.e., *imposters*. We opened the study to people who we knew were 30 years of age or younger, but we explicitly advertised the study as being only for people over age 50.

Method

Participants and Design

We set up the study on CloudResearch (Litman et al., 2017) and collected data from 100 participants within the U.S. Using MTurk Worker ID's, we opened the study to people whom we knew were 30 years of age or younger because they had previously indicated that age in demographic questions asked by CloudResearch. Since we knew all participants were actually 30 years old or younger, but we advertised the HIT as being for workers over 50, we knew anyone who took the study was misrepresenting their age. We advertised the HIT on MTurk as “ONLY for workers who are OVER 50” years of age. The study had a 24% bounce rate—i.e., 24% of workers who previewed the HIT did not accept it—and took longer than is typical to gather 100 responses (18 hours). All stimuli were the same as in Study 1. Participants received 50 cents to complete the study which we estimated would take 5 minutes. All data were collected in May 2019.

Results

Instrument Performance by Age

As in Study 1, we first obtained the AVI-S score by calculating a difference score, subtracting the percentage of correct contemporary items from the percentage of correct historical items. For the impostors assessed in this study, the correlation between self-reported age and score on the AVI-S was not significant ($r = -.162, p = .11$). As can be seen in Figure 2,

all age groups had negative mean scores, indicating that they performed better on the contemporary items compared to the historical items ($M = -26.07$, $SD = 44.62$; see Table S6 for group descriptives). Further, most participants were categorized by the instrument as young (receiving a score of 0 or lower), regardless of whether they claimed to be in their 20s (86.4%), 40s (100%), 50s (81.6%), 60s (70.6%) or 70s and over (100%). Thus, the AVI-S verified these participants as imposters. Table 2 presents a summary of these results.

Figure 2

Study 2 AVI-S Difference Scores by (Reported) Age Group

Note. AVI-S = Age Verification Instrument (Short). Error bars represent 95% confidence intervals. Age group reflects the age participants reported.

Discussion

In Study 2, our goal was to determine whether our instrument could “catch” people who misrepresented their age. By specifically inviting MTurk workers who we knew were under age 30 (Rosenzweig et al., 2016), while also clearly stating in the survey qualifications that the study was only for participants over age 50, we presented participants with an opportunity to lie about their age. Many of them took this opportunity. Indeed, though approximately one fifth of

participants reported their actual age, 72% claimed to be adults 50 or older. Their responses to the instrument, however, revealed their young age, with few participants (17% of the entire sample and 19.4% of participants claiming to be 50 or over) obtaining a positive difference score. Therefore, these results suggest the AVI is a powerful tool to detect imposters—catching over 80% of people who claimed to be older than they were to participate in a brief research opportunity.

Because the instrument relies on particular era-based knowledge of popular songs and TV shows, it may be less accurate in determining the age of populations that vary in era-based knowledge. Thus, in Study 3 we examine the cultural validity of the AVI with samples that differ by race, education level, and immigration status.

Study 3: Exploring Cultural Validity

After validating the AVI and confirming that it differentiates between younger and older participants, we turn to the instrument's cultural validity. Specifically, we test whether the instrument can distinguish between younger and older participants who are African American (Study 3a), whose education level does not exceed high school (Study 3b), and who immigrated to the U.S. (Study 3c). We created the AVI from popular songs and TV shows that we expected would transcend cultural differences in knowledge, given their wide popularity and mass appeal. Still, there are known differences in preferences for specific genres across race, ethnicity, and education level, among other factors (Mizell et al., 2005). Thus, we ran Study 3 to test the robustness of the AVI.

Method

Participants and Design

All data were collected in June and July 2019. Studies 3a and 3b were run on CloudResearch's MTurk Toolkit. Study 3c attempted to recruit very low-incidence groups (e.g., immigrants between the ages of 60-70) which cannot be easily recruited from MTurk. For this reason, we used CloudResearch's Prime Panels, which has a much larger participant pool that is more suitable for low-incidence studies (see Chandler et al., 2019) to obtain that sample. In each study, our samples targeted a total of 300 participants from six age groups (20s through 70s).

In Studies 3a and 3b, conducted on MTurk, we recruited 263 African Americans (Study 3a) and 264 adults with formal education equivalent to a high school diploma or less (Study 3b). Because MTurk's participant pool skews younger (Litman et al., 2020b) and we were targeting participants within already hard-to-reach groups, finding people at the top of our age range was difficult. After initially offering people 50 cents for a 5-minute study, we increased pay to \$1.50 and sent recruitment emails to all eligible participants in the CloudResearch database. Even with these adjustments, finding people over age 60 who fit the demographic requirements was difficult. Therefore, to compensate, we oversampled people aged 50-60 and closed both studies after 7 days.

Similarly, in Study 3c conducted on Prime Panels, we had trouble recruiting recent immigrants to the U.S. who were over age 60. Thus, we ended up with a sample of 264 adults and oversampled people in their 50s. Data collection ended after 14 days.

Procedure

As in Studies 1 and 2, we presented participants with the 19-item AVI. See Study 1 for full materials and procedure.

Results

Instrument Performance by Age

As in Studies 1 and 2, we first obtained the AVI-S score by calculating a difference score, subtracting the percentage of correct contemporary items from the percentage of correct historical items. Across the three samples, there were very few participants over the age of 70 (two in Study 3a, five in Study 3b, and eight in Study 3c), so we combined their data with the younger decade group (ages 60-70). The correspondence between age groups based on self-reported age and age in the CloudResearch database was high across studies 3a ($r = .96$) and 3b ($r = .99$), both $ps < .001$.

We ran a linear regression analysis predicting self-reported age with the AVI-S. In Studies 3a, 3b, and 3c, respectively, the AVI-S significantly predicted self-reported age ($r = .79$, $p < .001$; $r = .78$, $p < .001$; $r = .57$, $p < .001$) and explained 63%, 60%, and 32% of the variance in age.

There were significant differences between the group means in Study 3a ($F(4, 258) = 106.11$, $p < .001$), Study 3b ($F(4, 259) = 98.85$, $p < .001$), and Study 3c ($F(4, 259) = 35.49$, $p < .001$). In the African American (3a) and low education (3b) samples, the scores of all age groups differed from each other (all $p < .001$) except the two highest groups ($p = .902$ and $p = .823$, respectively). In the immigrant sample (3c), the two youngest groups did not differ from each other ($p = .99$) but did differ from all other groups (20s vs. 30s: $p = .023$, 30s vs. 40s: $p = .044$; all other comparisons: $p < .001$). Participants in their 40s also differed from all people in their 50s and 60s+ ($ps < .001$). The two oldest groups (50s and 60s+) did not differ from each other ($p = .58$). See Figure 3 for mean differences and see Tables S8-S10 for group descriptives.

As in Studies 1 and 2, we also tested whether self-reported age was more strongly correlated with the difference score compared to just the contemporary or historical items. For studies 3a, 3b, and 3c, self-reported age correlated less strongly with the contemporary items ($r =$

-.48, $r = -.43$, and $r = -.30$), than with the historical items ($r = .78$, $r = .79$, and $r = .49$), or with the difference score ($r = .79$, $r = .78$, and $r = .57$), respectively (all $ps < .001$).

As in Study 1, we also tested how accurate the AVI-S was at categorizing people into their age groups. On average, the instrument successfully categorized people across all three studies 85% of the time. In Study 3a, of the participants in their 50s and 60s or over, 88.8% and 92.9%, respectively, were correctly categorized as “older”. Participants in their 40s were split between the “older” (62.5%) and “younger” (37.5%) groups, but there were more in the “older” group. Of the participants in their 30s and 20s, 71.4% and 96.64%, respectively, were correctly categorized as “younger”. In Study 3b, of the participants in their 50s and 60s or over, 90.9% and 86.4%, respectively, were correctly categorized as “older”. Participants in their 40s were more evenly split between the “older” (51.1%) and “younger” (48.9%) groups. Of the participants in their 30s and 20s, 81.8% and 98.1%, respectively, were correctly categorized as “younger”. In Study 3c, of the participants in their 50s and 60s or over, 70.1% and 82.4%, respectively, were correctly categorized as “older”. Participants in their 40s were split between the “older” (29.5%) and “younger” (70.5%) groups, but there were more in the “younger” group. Of the participants in their 30s and 20s, 81.1% and 83.7%, respectively, were correctly categorized as “younger”. See Table 2 for a summary of these results.

Figure 3

Study 3 AVI-S Difference Scores by Age Group

Note. AVI-S = Age Verification Instrument (Short). Error bars represent 95% confidence intervals.

Discussion

In this final study, we sought to validate the AVI in samples that we had a priori reasons to believe might lack the era-based knowledge typical of the more general samples we used in Studies 1 and 2. Specifically, we examined whether the instrument differentiated between age groups in samples of African Americans, participants whose level of education did not exceed high school, and immigrants. Across all three samples, there was a significant linear effect of self-reported age on instrument performance, indicating that younger participants were better at

responding to the contemporary items and older participants were better at responding to the historical items.

Among the African American and low education samples, the results replicated the pattern we observed in Study 1—the two youngest groups had very negative scores, the two oldest groups had very positive scores, and the middle groups were close to a score of 0. In the immigrant sample, the same pattern emerged, but the differences between groups were less stark: the younger groups' scores were less negative, and the older groups' scores less positive. Further, the middle-aged group had slightly negative scores, indicating their knowledge was more similar to that of younger groups in the other samples. Although we can only speculate as to why this may be, one possibility is that the American culture that reached people's home countries was delayed by a few years, as was often the case before the internet. Regardless, the results of all three samples demonstrate the utility of the instrument for verifying participant age, even in different cultural groups.

General Discussion

People in online studies sometimes misrepresent themselves. Sometimes they do this out of mischievousness (Robinson-Cimpian, J.P, 2014) and sometimes they do it to qualify for studies they would otherwise be ineligible to participate in. Regardless of motive, participant misrepresentation threatens the validity of research. In the studies reported here, we proposed and tested a way to verify the age of online respondents: a test of era-based knowledge. The measure we created discriminated between people of different ages recruited from common online platforms, picked out imposters, and showed evidence of validity in studies with groups of different racial, educational, and cultural composition within the United States. Across all

studies, we found that the relative difference in knowledge between younger and older adults correctly categorized people by self-reported age between 74% and 98% of the time.

While our instrument consistently demonstrated a linear effect of age, the effects in Studies 3a, 3b, and 3c were smaller than those in Studies 1a and 1b. One explanation for this is that the effect among people with different cultural backgrounds is smaller (but still robust). Another possibility is that we did not recruit a large enough sample of adults over age 60. In all three studies exploring cultural validity, recruitment of people over age 60 was slow, so we oversampled people ages 50-60. Because adults in their 60s and 70s tend to produce some of the largest difference scores on our measure it is possible the results of Study 3 would be stronger with more older adults. Nevertheless, our instrument's ability to distinguish between older and younger adults even in this smaller range of ages and across different subcultures speaks to the robustness of the instrument.

A concern researchers may have in implementing an instrument like ours is that online studies provide easy access to internet searches, and thereby may threaten the validity of knowledge-based assessments. Our findings suggest this did not occur. Since young participants performed relatively poorly when responding to older items, and older participants performed relatively poorly when responding to younger items, it is unlikely people used the internet to look up answers. Indeed, and somewhat ironically, even in Study 2, when our sample consisted of people who were lying about their age, we saw little evidence of cheating (i.e., looking up answers). This may be because we explicitly told participants that performance would not affect compensation, because participants did not care to spend time looking up the answers, or for other reasons. Either way, evidence of cheating in our studies was low. Further, there are tools

that allow researchers to monitor when participants leave a task (see Permut et al., 2019) and indirect forms of data that may give insight into cheating (Steger et al., 2020).

At its core, our instrument was intended to address a contradiction with online research. The internet offers relatively easy access to participants, including older adults and many hard-to-reach samples. But at the same time, the internet makes it easy for people to misrepresent themselves and difficult for researchers to verify the identities of participants with whom there is little to no contact. Instruments like ours may allow researchers to be more confident they have successfully recruited their target group, thereby bolstering the validity and reliability of their findings. In the space below, we describe how instruments like ours may be used to verify the identity of online research participants and discuss some challenges to widespread implementation of such measures.

Applications

As is true of most scales, the longer they are, the better they capture what they are intended to measure. We reported the results of a shortened version of the AVI in the main text, and the results of the full version in the supplemental materials. In most cases, the longer scale's improvement was insubstantial, and for most uses, the benefits of a short scale outweigh the benefits of a longer scale that is just slightly better. However, researchers should be aware of the tradeoff in length vs. precision and use the expanded scale if greater precision is desired.

Researchers should note that the AVI is not particularly well-suited to pinpoint participants' exact age, since any young person could happen to be a fan of oldies and any older person could happen to be hip (and presumably would know not to call herself 'hip'). Although the percentage of such people was shown to be small in our studies, we would not recommend rejecting participants solely based on their performance on the instrument. Nonetheless, the AVI

provides a useful check, that, in combination with other measures, can be used to provide assurance with respect to participant age. In this respect, the AVI is like many data quality checks—one can never be certain if any particular participant is paying attention and answering earnestly. For example, when deciding to exclude outliers or participants who straight-line, the information from one measure is often ambiguous—the participant might truly strongly agree with all statements, or she might simply be clicking through. Researchers must make decisions (preferably prior to data collection) about how strictly they will apply exclusion criteria. If participant age is important to the study, researchers may choose to apply strict criteria and exclude anyone who receives an AVI score that does not align with their stated age, though this may result in some “false positive” cases in which honest participants are excluded.

Further, the instrument can be used to assess age at the group level. If one were attempting to recruit a sample of adults in their 60s, for example, and more than a handful of participants had negative scores, this should set off alarm bells and prompt researchers to look more closely at the rest of the participants’ data. While any participant may have idiosyncratic knowledge, on average, most of the older participants should have relatively more knowledge of older culture, and most of the younger participants should have relatively more knowledge of younger culture. At the group level, it is easy to determine if one has successfully recruited a group of older adults, or if the proportion of imposters in the sample is high. This type of group-level metric may be particularly useful when recruiting from a novel or open source of online participants and when additional confidence in the accuracy of participants’ demographics may be of value. If participants are not performing as expected, it might be prudent to run analyses with and without the suspicious participants (reporting both values, of course).

Challenges

While short screening measures that verify the demographic information of online respondents hold great potential, there are several challenges to widespread implementation of such measures. One such challenge is that the purpose needs to remain concealed from participants. If participants knew a vetting measure was intended to assess whether they actually belong to a demographic group, they would likely change their behavior. Our participants showed clear evidence of era-based knowledge because they did not know the purpose of our instrument. Although it may be hard to fake answers to factual questions (unless one looks up the answers online), past research demonstrates that imposters in online research often try to theorize about the people they are imitating. In trying to answer questions as they think X person would, imposters often show an exaggerated pattern of responses (e.g., Wessling et al., 2017). People under the age of 50, for example, report more interest in dietary fiber than people over age 50 report (Wessling et al., 2017). Therefore, it is important that the purpose of verification instruments remains concealed from participants. This may be more challenging if the main study is unrelated to the verification instrument.

A second challenge to widespread implementation of measures like the one we have developed here is that participants may become over-exposed to such measures with time. Just as experienced online participants have gradually become familiar with various attention checks (e.g., Hauser & Schwarz, 2016), extended use of any particular verification instrument would likely diminish its efficacy over time. For this reason, we recommend routinely updating the items and alternating between different verification measures. Although this approach requires more time and effort than simply reusing the same instrument, it will pay off in better data quality over time.

Finally, although in Study 3 we replicated our findings among groups with diverse ethnic and cultural backgrounds, the effects were smaller than previous studies—the percentages of correct categorizations were in the 70s and 80s, as opposed to the 90s as in previous studies. In other words, although the instrument still worked, it was not as robust as in Study 1 in which participants were not targeted based on race, education, or immigration status. While we expect it would be feasible to find items that do not differ by culture, it might be more prudent in some cases to develop multiple instruments for different groups (for example, creating separate instruments for White and Black Americans).

Future Directions

As mentioned above, we see the promise of the AVI in our general method more than in the specific items we used. Future age verification research should replicate the method using different items, and attempt to improve the accuracy of items for groups with different cultural backgrounds. Further, just as our verification method is not limited to the particular items used in the present study, it is also not limited to age. Perhaps one of the most fruitful avenues for future work is in verifying people's identities in studies that recruit hard-to-reach participants and are often extremely expensive to conduct. For example, researchers sometimes seek to recruit highly compensated professionals (e.g., CEOs, IT professionals) or people in very small demographic groups (e.g., parents of children with autism spectrum disorder, veterans). In some of these cases, it can cost more than 80x as much to recruit one of these participants compared to a typical participant². Asking participants in these samples to respond to a short quiz assessing their knowledge is a promising way to ensure precious resources do not go to waste.

² Based on quotes provided to us by panel providers

While we have focused on knowledge to verify participant identity, there is no reason to limit instruments to this narrow scope. Instruments could take a more applied focus, by assessing specific skills (e.g., if researchers are recruiting people who work at particular jobs, see Danilova et al., 2021). Additionally, future studies could attempt a cross-validation measure that incorporates several pieces of information to indicate the likelihood that someone is of a certain identity.

Conclusion

Online studies offer many advantages over more traditional methods of data collection. Researchers can collect more data from more diverse people in less time than they would otherwise be able to do. However, with these advantages come some drawbacks, one of which is participants' ability to misrepresent themselves. We created and validated a measure to assess participants' age group and determine if they are who they say they are. The measure correctly identified participants' age group based on the percentage of correct answers to items that reflect temporal cultural trends. Beyond the specific items we have used here, the verification method is a promising approach researchers can and should consider to reduce the risk of collecting faulty data.

Open Practices Statement

Data, analysis script, and supplemental materials are available on OSF:

https://osf.io/bn4xy/?view_only=7252e963f3bd4c0c981eed6ddd085ee8.

None of the studies were preregistered.

References

- Antoun, C., Zhang, C., Conrad, F. G., & Schober, M. F. (2016). Comparisons of online recruitment strategies for convenience samples: Craigslist, Google AdWords, Facebook, and Amazon Mechanical Turk. *Field Methods*, 28, 231-246.
doi:10.1177/1525822X15603149
- Beier, M. E., & Ackerman, P. L. (2001). Current-events knowledge in adults: An investigation of age, intelligence, and nonability determinants. *Psychology and Aging*, 16(4), 615–628.
<https://doi.org/10.1037//0882-7974.16.4.615>
- Burnham, M. J., Le, Y. K., & Piedmont, R. L. (2018). Who is Mturk? Personal characteristics and sample consistency of these online workers. *Mental Health, Religion & Culture*, 21(9-10), 934-944.
- Chandler, J., & Paolacci, G. (2017). Lie for a dime: When most prescreening responses are honest but most study participants are impostors. *Social Psychological and Personality Science*, 8, 500-508. <https://doi.org/10.1177%2F1948550617698203>
- Chandler, J., Rosenzweig, C., Moss, A. J., Robinson, J., & Litman, L. (2019). Online panels in social science research: Expanding sampling methods beyond mechanical turk. *Behavior Research Methods*, 51(5), 2022-2038. <https://doi.org/10.3758/s13428-019-01273-7>
- Chandler, J., Sisso, I., & Shapiro, D. (2020). Participant carelessness and fraud: Consequences for clinical research and potential solutions. *Journal of Abnormal Psychology*, 129(1), 49.
- Danilova, A., Naiakshina, A., Horstmann, S., & Smith, M. (2021, May). Do you really code? Designing and Evaluating Screening Questions for Online Surveys with Programmers. In 2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE) (pp. 537-548). IEEE.

- Fan, X., Miller, B. C., Christensen, M., Park, K. E., Grotevant, H. D., van Dulmen, M., ... & Bayley, B. (2002). Questionnaire and interview inconsistencies exaggerated differences between adopted and non-adopted adolescents in a national sample. *Adoption Quarterly*, 6, 7-27. https://doi.org/10.1300/J145v06n02_02
- Gadiraju, U., Kawase, R., Dietze, S., Demartini, G., (2015). Understanding malicious behavior in crowdsourcing platforms: the case of online surveys. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI 2015, Seoul, Republic of Korea, April 18-23, pp. 1631e1640. <http://dx.doi.org/10.1145/2702123.2702443>
- Hauser, D. J., & Schwarz, N. (2016). Attentive turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods*, 48(1), 400-407. doi:<https://doi.org/10.3758/s13428-015-0578-z>
- Janssen, S. M. J., Murre, J. M. J., & Meeter, M. (2008). Reminiscence bump in memory for public events. *European Journal of Cognitive Psychology*, 20(4), 738–764. <https://doi.org/10.1080/09541440701554409>
- Kahana, M. J., Howard, M. W., Zaromb, F., & Wingfield, A. (2002). Age dissociates recency and lag recency effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(3), 530.
- Kan, I. P., & Drummy, A. B. (2018). Do imposters threaten data quality? An examination of worker misrepresentation and downstream consequences in Amazon’s Mechanical Turk workforce. *Computers in Human Behavior*, 83, 243-253.
- Kramer, J., Rubin, A., Coster, W., Helmuth, E., Hermos, J., Rosenbloom, D., . . . Liljenquist, K. (2014). Strategies to address participant misrepresentation for eligibility in Web-based

- research. *International Journal of Methods in Psychiatric Research*, 23(1), 120-129.
doi:10.1002/mpr.1415
- Litman, L. & Robinson, J. (2020). *Conducting Online Research on Amazon Mechanical Turk and Beyond*. Sage Academic Publishing. Thousand Oaks: CA
- Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior research methods*, 49, 433-442.
- Litman, L., Robinson, J., & Rosenzweig, C. (2020a). Beyond Mechanical Turk: Using online market research platforms. In L. Litman and J. Robinson (Eds.) *Conducting Online Research on Amazon Mechanical Turk and Beyond* (217-233). Sage Academic Publishing. Thousand Oaks: CA
- Litman, L., Robinson, J., & Rosenzweig, C. (2020b). Sampling Mechanical Turk workers: Problems and solutions. In L. Litman and J. Robinson (Eds.) *Conducting Online Research on Amazon Mechanical Turk and Beyond* (148-172). Sage Academic Publishing. Thousand Oaks: CA
- MacInnis, C. C., Boss, H. C., & Bourdage, J. S. (2020). More evidence of participant misrepresentation on Mturk and investigating who misrepresents. *Personality and Individual Differences*, 152, 109603. <https://doi.org/10.1016/j.paid.2019.109603>
- Mizell, L. (2005). *Music Preferences in the US: 1982-2002*. National Endowment for the Arts.
- Permut, S., Fisher, M., & Oppenheimer, D. M. (2019). Taskmaster: A tool for determining when subjects are on task. *Advances in Methods and Practices in Psychological Science*, 2(2), 188-196. <https://doi.org/10.1177/2515245919838479>

Pew Research Center (2021, April 7). Internet/Broadband Fact Sheet.

<https://www.pewresearch.org/internet/fact-sheet/internet-broadband/>

Robinson-Cimpian, J. P. (2014). Inaccurate estimation of disparities due to mischievous responders: Several suggestions to assess conclusions. *Educational Researcher*, 43(4), 171-185.

Robinson, J., Litman, L., & Rosenzweig, C. (2020). Who are the Mechanical Turk workers? In L. Litman and J. Robinson (Eds.) *Conducting Online Research on Amazon Mechanical Turk and Beyond* (121-147). Sage Academic Publishing. Thousand Oaks: CA

Rosenzweig, C., Robinson, J., & Litman, L. (2016, January) Are They Who They Say They Are?: Reliability and Validity of Web-Based Participants' Self-Reported Demographic Information. Poster presented at the 18th Society for Personality and Social Psychology Annual Convention, San Antonio, TX.

Shatz, I. (2017). Fast, free, and targeted: Reddit as a source for recruiting participants online. *Social Science Computer Review*, 35(4), 537-549.

Siegel, J. T., & Navarro, M. N. (2019). A conceptual replication examining the risk of overtly listing eligibility criteria on Amazon's Mechanical Turk. *Journal of Applied Social Psychology*, 49, 239-248. doi: 10.1111/jasp.12580

Siegel, J. T., Navarro, M. A., Thomson, A. L. (2015). The impact of overtly listing eligibility requirements on MTurk: An investigation involving organ donation, recruitment scripts, and feelings of elevation. *Social Science and Medicine*, 142, 256-260

Steger, D., Schroeders, U., & Wilhelm, O. (2020). Caught in the act: Predicting cheating in unproctored knowledge assessment. *Assessment*, 1–14.
<https://doi.org/10.1177/1073191120914970>

Wessling, K., Huber, J., & Netzer, O. (2017). MTurk character misrepresentation: Assessment and solutions. *Journal of Consumer Research*, 44, 211-230.

<https://doi.org/10.1093/jcr/ucx053>

Zimprich, D. (2020). Individual differences in the reminiscence bump of very long-term memory for popular songs in old age: A non-linear mixed model approach. *Psychology of Music*, 48(4), 547–563. <https://doi.org/10.1177/0305735618812199>

Author Contributions

	RH	AJM	IR	NB	CR	JR	LL
Conceptualization							
Data curation							
Formal analysis							
Investigation							
Methodology							
Project administration							
Software							
Writing – original draft							
Writing – review & editing							