Encouraging self-blinding in hiring

Sean Fath, Richard P. Larrick, & Jack B. Soll

Method & Analysis

More About Relevant Studies of Default Effects

Decision-makers are more likely to select options when they must opt out to avoid selection than when they must opt in to select them. Similar patterns of effects emerge when decision-makers are tasked with selecting specific options from an initial broad set by a process of inclusion (that is, selecting the best alternatives from the initial set) versus exclusion (that is, rejecting the worst alternatives from the initial set). Through a process of inclusion, any options not selected from the initial set are dropped from consideration by default; through a process of exclusion, any options not rejected from the initial set remain to be considered by default. Consistent with research on default effects, research on inclusion versus exclusion frames demonstrates that decision-makers tend to generate larger subsets of alternatives when using an exclusion frame (rejecting bad options) versus when they use an inclusion frame (including good options; Yaniv & Schul, 1997). In one study in a hiring context, participants selected more applicants to be interviewed out of a larger pool when they were using an exclusion frame (rejecting those who should not be interviewed) versus an inclusion frame (including those who should be interviewed; Huber et al., 1987).

Default effects have a powerful influence on behavior because people are likely to make choices that follow the path of least resistance, and the easiest thing to do is often the thing that requires no choice or action in the first place (that is, doing nothing at all; Choi et al., 2002). Similarly, one reason for differences in the size of choice sets generated under inclusion versus exclusion frames is that decision-makers use stricter criteria for retaining an option under inclusion than under exclusion (Huber et al., 1987)—the necessary attentiveness to and deliberation between different options is higher under an inclusion frame than under an exclusion frame (Levin et al., 2000). Notably, default effects can also produce suboptimal outcomes. For instance, while automatic enrollment in 401(k) plans can increase the number of employees who ultimately enroll in such plans, these employees often select default (low) rates of contribution, which may not match the employee's savings goals (Madrian & Shea, 2001).

Posttest: Additional Information

We explored whether the effects of our experimental manipulations on participants' likelihood of choosing a given set of information about the job candidate varied by whether the information was useful or biasing. First, it was necessary to confirm that the five items of information available to participants that we prejudged to be useful or diagnostic of job performance—name of college, major in college, previous work experience, job-related skills, and references—were indeed perceived by participants to be relatively more useful or diagnostic of job performance than biasing. Similarly, we sought to confirm that the two items of information available to participants that we prejudged to be biasing—either race and gender (*transparent bias* condition) or picture and name (*nontransparent bias* condition)—were indeed perceived to be relatively more biasing than useful or diagnostic of job performance. That is, we could not assess differences resulting from our experimental manipulation on information choice as a function of item status as useful or biasing if participants did not perceive the items to be differently useful or biasing in the manner we expected.

We recruited 104 participants with experience making hiring decisions in their careers (a sample that did not overlap with the sample that was used in the article) via Prolific Academic (*Mage* = 41.62 years, *SDage* = 11.63 years, 59.6% women). Participants had an average of 21.30 years of work experience and estimated having made an average of 27.14 hiring decisions in their careers. Each participant read, "Think about the typical manager making a hiring decision. On the next page, you will see multiple different pieces of information about job applicants that the typical manager making a hiring decision might have." Next, participants rated the nine items of information listed above (five useful and four biasing), in a randomized order, in terms of the degree to which they were (a) relevant or useful and (b) biasing on a scale of 1 (*not at all*) to 7 (*very*). The order of these two questions was itself randomized. The question about usefulness perceptions read, "To what extent would you judge the pieces of information listed below as relevant to, or useful for, the hiring decision of a typical manager?" The question about potential for bias read, "To what extent would you judge the pieces of information listed below potentially biasing for the hiring decision of a typical manager?"

We expected that, on average, participants would judge the five items we prejudged to be useful to be relatively more useful or relevant than biasing, whereas participants would judge the four items we prejudged to be biasing to be relatively more biasing than useful. To assess this, we subtracted participants' ratings of each item's degree of potential for bias from their ratings of usefulness. Table S1 below details these net ratings. As can be seen in Table S1, four of the items we prejudged to be useful were indeed rated by participants as being relatively more useful than biasing—major in college, previous work experience, job-related skills, and references. Each of the items we prejudged to be biasing were rated by participants as relatively more biasing than useful. However, contrary to our expectations, one of the items we prejudged to be useful—name of the job applicant's college—was actually rated slightly more biasing than useful to a hiring decision by our participants with hiring experience. This

Table S1. Posttest ratings of item usefulness versus potential for bias

Item	Useful/relevant		Biasing		Net rating (useful – biasing)	
	M	*SD*	*M*	*SD*	*M*	*SD*
Name of college	4.05	1.51	4.09	1.76	–0.04	1.98
Major in college	5.40	1.32	3.92	1.75	1.48	1.96
Previous work exp.	6.53	0.79	4.28	2.05	2.25	2.19
Job-related skills	6.76	0.55	4.17	2.21	2.59	2.31
References	5.60	1.37	4.04	1.83	1.56	2.16
Race	1.81	1.23	4.95	2.21	–3.14	2.42
Gender	2.20	1.42	4.74	2.05	–2.54	2.23
Picture	2.23	1.56	5.03	1.97	–2.80	2.53
Name	2.70	1.96	4.04	1.96	–1.34	2.77

Note. Previous work exp. = previous work experience.

helps to explain why the "name of college" item was generally selected by a lower percentage of participants than each of the other useful items (see Table S2). For instance, in the opt-in conditions, between 79% and 96.6% of participants selected the other four useful items. Conversely, the name of college item was selected by roughly 58% of participants in the opt-in conditions. Given that the name of college item did not appear to be perceived by participants in the main study in the manner we expected—as a useful piece of information, diagnostic of job performance—we exclude it from the following analyses, contrasting the effects of experimental condition on participants' choices to see the useful versus biasing information.

Method, Analyses, & Results: Additional Information

Our study had a 2 (default: opt in versus opt out) × 2 (subject: self versus other) × 2 (bias: transparent versus nontransparent) between-subjects design. Participants were randomly assigned to receive no information by default and to opt in to any information they wished to receive (*opt-in condition*) or to receive all information by default and to opt out of any information they wished to avoid (*opt-out condition*). Participants were randomly assigned to make a choice about which information items to receive, were they personally making the hiring decision (*self condition*), or which they wished to assign to someone else, were someone else making the decision (*other condition*). Finally, participants were randomly assigned to have either two transparently biasing items of information available to them in the full suite of items (*transparent bias condition*) or two nontransparently biasing items available to them (*nontransparent bias condition*). Only participants who provided complete survey responses were included in analyses.

Choice to View Biasing Information

The percentage of participants choosing to see each item, organized by condition, is displayed in Table S2. First, we explored the effects of the above manipulations on participants' likelihood of choosing the biasing information about the mock job candidate. We used linear probability models for these analyses (Angrist & Pischke, 2009; R Core Team, 2021). In a linear probability model, a binary dependent variable is regressed directly on the independent variables. The resulting coefficients are the same as those in ordinary least squares, except that the standard errors are adjusted to account for heteroscedasticity. In our case, because each participant provided multiple observations (that is, observations are clustered within participant), we used cluster-robust standard errors to determine statistical significance (Angrist & Pischke, 2009). This procedure adjusts the standard errors to account for both heteroscedasticity and correlated observations. We implemented this in R using the sandwich (Zeileis et al., 2020) and miceadds (Robitzsch & Grund, 2021) packages.

Table S2. Percentage of participants who chose to view items by condition

Items	All items provided by default (opt out)		No items provided by default (opt in)	
	Choice for self	Choice for someone else	Choice for self	Choice for someone else
Useful				
Major in college	86.6	91.7	81.9	83.9
Prev. work exp.	93.8	97.1	96.1	96.1
Job-related skills	94.7	97.1	96.6	96.6
References	90.4	88.3	84.8	79.0
Name of college	75.6	78.2	58.3	58.0
Biasing				
Race	17.6	14.7	19.0	5.0
Gender	30.4	22.0	18.0	9.0
Picture	33.6	21.6	26.0	14.3
Name	66.4	50.5	50.0	35.2

Note. Prev. work exp. = previous work experience.

Table S3. Choice to view the biasing information

Variables and interaction terms	*b*	*SE*	*t*
Intercept	.267***	.013	21.06
Bias	.103***	.013	8.14
Default	.046***	.013	3.62
Subject	–.055***	.013	–4.32
Bias × Default	.009	.013	0.68
Bias × Subject	–.011	.013	–0.86
Default × Subject	.008	.013	0.61
Bias × Default × Subject	–.013	.013	–1.00

****p* < .001.

Using the above procedure, we regressed participants' choice to view the biasing information (1 = yes, 0 = no) on default (1 = opt out, –1 = opt in), subject (1 = other, –1 = self), bias (1 = nontransparent, –1 = transparent), and the interaction terms. Results are reported below in Table S3. There were no significant interactions, but we found significant main effects of each factor in our experiment.

Preference for Useful Versus Biasing Information

We explored whether the default (opt in versus opt out) and subject (self versus other) manipulations had different effects on participants' likelihood of choosing the useful versus biasing information. Because the type of biasing information varied between participants as a function of experimental condition (transparent versus nontransparent bias), we first contrasted participants' likelihood of choosing the useful versus biasing information in the transparent bias condition. As above, we used linear probability models, clustering standard errors within participant, with cluster-robust standard errors. We regressed participants' choice to view the information (1 = yes, 0 = no) on default (1 = opt out, –1 = opt in), subject (1 = other, –1 = self), information (1 = useful, –1 = biasing), and the interaction terms. Results are reported in Table S4. The significant Information × Subject interaction indicates that participants were more likely to choose the (transparently) biasing information for themselves than assign it to someone else, but this difference was attenuated for the useful information.

Table S4. Choice to view useful versus biasing information

Variables and interaction terms	*b*	*SE*	*t*
Intercept	.543***	.009	61.62
Information	.379***	.010	37.59
Default	.022*	.009	2.54
Subject	–.017*	.009	–1.96
Information × Default	–.015	.010	–1.47
Information × Subject	.027**	.010	2.63
Default × Subject	.015	.009	1.70
Information × Default × Subject	–.005	.010	–0.527

Note. Participants were assigned to the transparent bias condition.
*p < .05. **p < .01. ***p < .001.

Next, we contrasted participants' likelihood of choosing the useful versus biasing information in the nontransparent bias condition. As in previous analyses, we used linear probability models, clustering standard errors within participant, with cluster-robust standard errors. We regressed participants' choice to view the information (1 = yes, 0 = no) on default (1 = opt out, –1 = opt in), subject (1 = other, –1 = self), information (1 = useful, –1 = biasing), and the interaction terms. Results are reported in Table S5. The significant Information × Subject interaction indicates that participants were more likely to choose the (nontransparently) biasing information for themselves than to assign it to someone else, but this difference was attenuated for the useful information. The significant Information × Default interaction indicates that participants were more likely to ultimately choose the (nontransparently) biasing information when they were automatically provided with all information than when they were blind to all information by default, but this default effect was attenuated for choices to receive the useful information.

Deliberate Ignorance

Broadly, participants in our study avoided receiving biasing information for themselves or providing it to others. This is consistent with the concept of deliberate ignorance, in which an actor proactively avoids receiving certain information (Hertwig & Engel, 2016). There are many reasons for one to engage in deliberate ignorance, such as to avoid psychologically traumatic knowledge (for example, whether one will develop a debilitating disease; Yaniv et al., 2004), to avoid information that would challenge a core belief (for example, whether a favored politician is corrupt; Sweeny et al., 2010), and to avoid information that could foster bias (for example, whether a job candidate is from a racial minority group; Hertwig & Engel, 2016). In our study, participants appeared to

Table S5. Choice to view useful versus biasing information

Variables and interaction terms	*b*	*SE*	*t*
Intercept	.643***	.010	62.87
Information	.273***	.011	25.47
Default	.027**	.010	2.69
Subject	–.034***	.010	–3.31
Information × Default	–.027*	.011	–2.53
Information × Subject	.032**	.011	2.97
Default × Subject	.0004	.010	0.04
Information × Default × Subject	.005	.011	0.50

Note. Participants were assigned to the nontransparent bias condition.
*p < .05. **p < .01. ***p < .001.

favor self-blinding (that is, deliberate ignorance) for this latter reason—as a means to avoid bias in judgments.

Limitations

Of course, while we argue that participants in our study adopted self-blinding as a means to avoid bias, it is possible that, instead, participants avoided receiving biasing information to appear less biased. On the one hand, this alternative explanation—that participants' responses were driven by experimental demand—is endemic to nearly all research of the type we conducted: a hypothetical experiment performed online. We leveraged this medium to carefully test our hypotheses in a controlled, rigorous manner. On the other hand, other research demonstrates that incentives to take an evaluative task seriously—which should suppress demand-driven responses—do not affect self-blinding preferences (Fath et al., 2022). Still, it would be useful for future researchers to explore whether self-blinding decisions change in situations where decision-makers are versus are not accountable to others for their decisions.

Other Routes to Achieving Diversity Hiring Goals

We note that self-blinding initiatives, if adopted, should be used in conjunction with other diversity hiring solutions implemented at other stages of the hiring pipeline. Solutions applicable to the recruiting stage of the hiring pipeline (before initial screens are conducted), include, but are not limited to, the establishment of pipeline programs at historically Black colleges and universities (Office of Federal Contract Compliance Programs, n.d.) and the reduction of gendered or racialized wording in job advertisements (Gaucher et al., 2011). Moreover, although nudges to encourage self-blinding in initial screening decisions may increase the number of members of marginalized groups who make it to the interview stage, different solutions will be required at the interview stage, as interviews are unlikely to be conducted blind. Multiple bias-reduction strategies can be implemented to reduce discrimination in interviewing, including using structured interviews (van der Zee et al., 2002), avoiding panels of interviewers that are homogenous in terms of race or gender (Basi, 2021; Prewett-Livingston et al., 1996), and moving away from notions of cultural fit as an assessment criterion (Rivera, 2020).

references

Angrist, J. D., & Pischke, J. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.

Basi, S. (2021, January 20). *Diverse interview panels: A win-win for candidates and companies*. LinkedIn. https://www.linkedin.com/pulse/diverse-interview-panels-win-win-candidates-companies-basi-ph-d-/

Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2002). Defined contribution pensions: Plan rules, participant choices, and the path of least resistance. *Tax Policy and the Economy*, *16*, 67–113. https://doi.org/10.1086/654750

Fath, S. B., Larrick, R. P., & Soll, J. (2022). Blinding curiosity: Exploring preferences for "blinding" one's own judgment. *Organizational Behavior and Human Decision Processes*, *170*, Article 104135. https://doi.org/10.1016/j.obhdp.2022.104135

Gaucher, D., Friesen, J., & Kay, A. C. (2011). Evidence that gendered wording in job advertisements exists and sustains gender inequality. *Journal of Personality and Social Psychology*, *101*(1), 109–128. https://doi.org/10.1037/a0022530

Hertwig, R., & Engel, C. (2016). Homo ignorans: Deliberately choosing not to know. *Perspectives on Psychological Science*, *11*(3), 359–372. https://doi.org/10.1177/1745691616635594

Huber, V. L., Neale, M. A., & Northcraft, G. B. (1987). Decision bias and personnel selection strategies. *Organizational Behavior and Human Decision Processes*, *40*(1), 136–147. https://doi.org/10.1016/0749-5978(87)90009-4

Levin, I. P., Huneke, M. E., & Jasper, J. D. (2000). Information processing at successive stages of decision making: Need for cognition and inclusion–exclusion effects. *Organizational Behavior and Human Decision Processes*, *82*(2), 171–193. https://doi.org/10.1006/obhd.2000.2881

Madrian, B. C., & Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *The Quarterly Journal of Economics*, *116*(4), 1149–1187. https://doi.org/10.1162/003355301753265543

Office of Federal Contract Compliance Programs. (n.d.). *HBCU recruitment strategies*. U.S. Department of Labor. https://www.dol.gov/agencies/ofccp/compliance-assistance/outreach/hbcu-initiative/recruitment-strategies

Prewett-Livingston, A. J., Feild, H. S., Veres, J. G., III, & Lewis, P. M. (1996). Effects of race on interview ratings in a situational panel interview. *Journal of Applied Psychology*, *81*(2), 178–186. https://doi.org/10.1037/0021-9010.81.2.178

R Core Team. (2021). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. https://www.R-project.org/

Rivera, L. (2020, August 4). *Stop hiring for "cultural fit."* Kellogg Insight. https://insight.kellogg.northwestern.edu/article/cultural-fit-discrimination

Robitzsch, A., & Grund, S. (2021). *Miceadds: Some additional multiple imputation functions, especially for 'mice.'* The Comprehensive R Archive Network. https://CRAN.R-project.org/package=miceadds

Sweeny, K., Melnyk, D., Miller, W., & Shepperd, J. A. (2010). Information avoidance: Who, what, when, and why. *Review of General Psychology*, *14*(4), 340–353. https://doi.org/10.1037/a0021288

van der Zee, K. I., Bakker, A. B., & Bakker, P. (2002). Why are structured interviews so rarely used in personnel selection? *Journal of Applied Psychology*, *87*(1), 176–184. https://doi.org/10.1037/0021-9010.87.1.176

Yaniv, I., Benador, D., & Sagi, M. (2004). On not wanting to know and not wanting to inform others: Choices regarding predictive genetic testing. *Risk, Decision and Policy*, *9*(4), 317–336. https://doi.org/10.1080/14664530490896573

Yaniv, I., & Schul, Y. (1997). Elimination and inclusion procedures in judgment. *Journal of Behavioral Decision Making*, *10*(3), 211–220. https://doi.org/10.1002/(SICI)1099-0771(199709)10:3<211::AID-BDM250>3.0.CO;2-J

Zeileis, A., Köll, S., & Graham, N. (2020). Various versatile variances: An object-oriented implementation of clustered covariances in R. *Journal of Statistical Software*, *95*(1), 1–36. https://doi.org/10.18637/jss.v095.i01