Cogn Affect Behav Neurosci (2013) 13:674–685
DOI 10.3758/s13415-013-0202-6

How to produce personality neuroscience research with high statistical power and low additional cost

Raymond A. Mar · R. Nathan Spreng · Colin G. DeYoung

Published online: 28 August 2013

Abstract Personality neuroscience involves examining relations between cognitive or behavioral variability and neural variables like brain structure and function. Such studies have uncovered a number of fascinating associations but require large samples, which are expensive to collect. Here, we propose a system that capitalizes on neuroimaging data commonly collected for separate purposes and combines it with new behavioral data to test novel hypotheses. Specifically, we suggest that groups of researchers compile a database of structural (i.e., anatomical) and resting-state functional scans produced for other task-based investigations and pair these data with contact information for the participants who contributed the data. This contact information can then be used to collect additional cognitive, behavioral, or individual-difference data that are then reassociated with the neuroimaging data for analysis. This would allow for novel hypotheses regarding brain–behavior relations to be tested on the basis of large sample sizes (with adequate statistical power) for low additional cost. This idea can be implemented at small scales at single institutions, among a group of collaborating researchers, or perhaps even within a single lab. It can also be implemented at a large scale across institutions, although doing so would entail a number of additional complications.

Keywords Personality · Individual differences · Neuroscience · Neuroinformatics · Sample size · Statistical power

R. A. Mar (✉)
Department of Psychology, York University, 4700 Keele St. W., Toronto, ON M5T2X5, Canada
e-mail: mar@yorku.ca

R. N. Spreng
Department of Human Development, Cornell University, Ithaca, NY, USA

C. G. DeYoung
Department of Psychology, University of Minnesota, Minneapolis, MN, USA

Small samples and resultant low statistical power are associated with a number of inferential problems in all domains of research, including neuroimaging. Studies with low power increase the likelihood of inflated effect size estimates within studies (Yarkoni, 2009; cf. Jennings & Van Horn, 2012), and using more stringent alpha thresholds in small samples merely lowers power further, thus exacerbating the problem of inflated effect sizes, while increasing the frequency of false negatives (Gonzalez-Castillo et al., 2012; Thyreau et al., 2012; Yarkoni & Braver, 2010). Small samples and inadequate statistical power also produce an increased proportion of false positives relative to true positives (Green et al., 2008, Box 1), increasing the likelihood that falsely positive "statistically significant" results will enter the literature, even after applying appropriate statistical thresholds (Button et al., 2013; Pashler & Harris, 2012). Attaining adequate statistical power is even more difficult when one is interested in individual differences, because of the need for larger sample sizes relative to comparisons of mean differences, so power issues are especially important to consider for personality neuroscience. In this article, we outline how low statistical power is particularly difficult to overcome for personality neuroscience investigations and propose a novel database approach that can help to produce personality neuroscience studies with large samples for low additional cost.

Personality neuroscience

Personality neuroscience entails the examination of how variability among individuals on cognitive, emotional, motivational, or behavioral dimensions (e.g., extraversion, intelligence, empathic ability) is related to neural variables. This approach has uncovered a number of interesting phenomena based on a variety of neural variables, including the size of brain structures, functional connectivity between brain regions, and white matter organization. For example, the size and structure of different brain regions have been found to correspond to individual differences in the Big Five personality traits (DeYoung et al., 2010; Hu et al., 2011; Kapogiannis, Sutin, Davatzikos, Costa, & Resnick, 2012), empathy and social cognition (Banissy, Kanai, Walsh, & Rees, 2012; Holmes et al., 2012; cf. Mills, Lalonde, Clasen, Giedd, & Blakemore, 2012), self-reported social network size (Bickart, Wright, Dautoff, Dickerson, & Barrett, 2011), online social network size (based on Web sites such as Facebook; Kanai, Bahrami, Roylance, & Rees, 2012), and even perceptual rivalry for ambiguous drawings (Kanai, Bahrami, & Rees, 2010). Behavioral correlates of individual differences in white matter integrity have also been discovered, with white matter differences predicting variability in the discounting of future rewards (Yu, 2012) and self-reported empathy (Parkinson & Wheatley, 2012), as well as cognitive performance and gender differences (DeCarli et al., 1995; Gur et al., 1999). With respect to resting state functional connectivity, patterns of connectivity for different regions have been found to differ between the sexes (Biswal et al., 2010; Kilpatrick, Zald, Pardo, & Cahill, 2006), correlate with reading performance (Koyama et al., 2011), and covary with self-reported inner thoughts (Andrews-Hanna, Reidler, Huang, & Buckner, 2010; Doucet et al., 2012) and the Big Five personality traits (Adelstein et al., 2011), to name a few examples. Personality neuroscience studies have generated a great deal of interest and often appear in high-impact journals. They are, however, difficult to produce, because properly examining interindividual variability in a reliable way requires large sample sizes.

Statistical power for personality neuroscience

There are a number of reasons why large samples are required to attain adequate statistical power for personality neuroscience studies. Not surprisingly, if you are interested in the variability that exists between individuals, examining few individuals makes it unlikely that you are capturing an acceptable portion of the variability of interest. Unless one takes steps to maximize the chances of sampling from the full range of the distribution—through preselection based on pretesting for example—small random samples of the population will be biased toward the mean and are unlikely to represent measurements at either tail of the distribution (assuming a normal distribution). Even after preselecting to achieve representativeness within the predictor, increased sampling error in small samples can lead to problematic outliers on the criterion variable. In other words, studies with small sample sizes often suffer from a restriction of range and from multivariate outliers, attenuating the possibility of observing associations and obscuring the true relationship across the full distribution of scores.

In personality neuroscience, correlational methods such as regression are employed to study associations between individual differences in brain variables and behavior. Unfortunately, correlations are especially vulnerable to the problems associated with small sample sizes. From a statistical standpoint, the uncertainty that surrounds any estimate of relationship between two variables is greatly influenced by the number of individuals examined. Correlations are much more susceptible to the presence of outliers than are group or condition means, and so they have larger confidence intervals given the same sample size. This can be seen in calculations of statistical power for correlations versus t-tests (Cohen, 1988). For example, in order to have 80 % power to detect a correlation of .25 using an alpha threshold of .05, one requires 123 participants. The same degree of power to detect an equivalent effect-size ($d = 0.52$) for a one-sample t-test, however, requires only 31 participants. (All calculations based on the PWR package in R created by Stéphane Champely.)

In order to illustrate the importance of sample size for correlation estimates, consider a correlation (r) of .25, which falls in the center of the middle third of all correlations observed for measures that do not share method (Hemphill, 2003). For a correlation of .25 and a sample size of 20 individuals, the 95 % confidence interval for this association falls between −.22 and .63.[1] In other words, with 20 participants, one should have very little confidence in whether a correlation of typical size is positive or negative or whether an association exists at all. Keep in mind that studies provide sample-based estimates of population-level effects, with the confidence we have in these estimates greatly affected by sample size. It is only when correlations are .45 or higher that the low end of a 95 % confidence interval falls above 0 (from .01 to .74) based on 20 participants. Unfortunately, effects equivalent to a correlation of .45 or higher have historically been quite rare within psychological and medical research (Hemphill, 2003; Meyer et al., 2001). In the case of individual-difference research, the average correlation is estimated to be about .24 (SD = .17, median = .21; Fraley & Marks, 2007). On the basis of this estimation, a correlation of .45 is about 1.2 standard deviations above the mean. Put another way, around 88 % of all correlations fall below this

[1] Following the detection of a significant correlation in one's own data, the robustness of this association may be evaluated with a bootstrap resampling procedure (Efron & Tibshirani, 1986, 1993; Lee & Rodgers, 1998). Bootstrapping permits the calculation of confidence intervals that better estimate population parameters, as compared with the Fisher's r to z transformation used in our example, by using one's own data as a basis for estimation. These confidence intervals can then be used to evaluate whether estimates of population correlations are likely to include 0. Bootstrapping procedures are incorporated into a number of common statistical packages, including SPSS (IBM Corp.) and MATLAB (MathWorks, Inc.).

correlation of .45. As a consequence, nearly 90 % of all true associations will not be detected when correlations with a sample size of 20 participants are conducted.

What sample size is needed in order to have adequate power for a personality neuroscience study? Let's assume that the phenomena we are interested in are of the same magnitude as those typically observed in psychological and medical research in general. In this case, the average effect size is equivalent to an r of around .25, or a d of 0.52 (Fraley & Marks, 2007; Hemphill, 2003; Meyer et al., 2001). For a correlation of this magnitude, you would need a minimum of 67 participants in order to have a 95 % confidence interval that ranges from .01 to .46. Of course, the larger the sample size, the better off an analysis will be. Statistically speaking, it would not be at all surprising for the effect size of a particular phenomenon to fall one half a standard deviation below the average effect-size (r = .17, d = 0.35; based on the standard deviations reported by Fraley & Marks, 2007). About 70 % of all effect-sizes are half a standard deviation below the mean or greater. Detecting a correlation of this magnitude requires a sample of at least 150 participants for a correlational analysis (based on a 95 % confidence interval that does not include 0). In our opinion, a sample size of 150 participants or greater is recommended for examining individual differences within neuroscience, since this affords 95 % confidence in detecting about 70 % of all effect sizes based on known estimates.

In sum, for personality neuroscience studies to be credible, they require large sample sizes of at least 150 participants or more, unless one holds to the assumption that the effects in personality neuroscience are much larger, on average, than what has previously been observed for medical or individual-difference research. This does not seem likely to be the case, however. For one thing, no compelling reason exists to think that effect sizes would be larger for personality neuroscience research, as compared with other forms of biological or individual-difference research. Additionally, simulation studies and mathematical analyses have shown that even if only four causal variables completely determine the value of some outcome variable, the correlation between any one of those four variables and the outcome will not be higher than about .45 (Ahadi & Diener, 1989; Strube, 1991). Given that the brain is a complex system and variation in many brain processes and parameters are likely to have an influence on any given phenotypic trait or behavior (Zuckerman, 2005), it is unlikely that as few as four neural variables fully determine any given trait or behavior. We should therefore expect important effects to be in a similar range of magnitudes in neuroimaging research as they are in behavioral research, with associations greater than .45 a very rare occurrence. Some may have the impression that larger correlation values are commonly observed within neuroimaging; this may be because effect-size estimates are inflated as a result of low statistical power (Button et al., 2013).

In sum, the problem of attaining adequate statistical power is especially pronounced for those interested in studying how the brain relates to interindividual variability in psychological traits. This is due both to the need to capture an adequate breadth of variability for the dimensions of interest and to the need for larger samples to achieve similar levels of power for correlation analyses, relative to mean comparisons. With large sample sizes typically comes a high financial cost, placing sufficiently powered personality neuroscience investigations out of reach for many researchers. In what follows, we outline an approach that allows active neuroimagers to produce personality neuroscience studies with large samples for little additional cost, by capitalizing upon neuroimaging data collected for other purposes.

Addressing the problem of power in personality neuroscience

Because the cost of collecting large samples is prohibitive for most researchers, the obvious solution is for groups of researchers to pool their data (Poldrack, 2012; Yarkoni, Poldrack, Van Essen, & Wager, 2010). Data sharing and data mining are clearly advantageous when it comes to the problem of sample size and statistical power. However, researchers are notoriously reticent to share their data, for a number of reasons. These include the concern that after sharing their data, another researcher might reanalyze these data and uncover a novel finding that the original researcher therefore misses out on reporting. Another concern is that after data are shared, some other researcher might reanalyze these data and find problems with the original analysis. Rather than try to combat these concerns, we propose a different solution that circumvents them. Neuroimagers routinely collect data in the course of task-based studies that they are unlikely to reanalyze to test new hypotheses, and these data are also unlikely to be useful for questioning published results. These include high-resolution anatomical scans, scans dedicated to the visualization of white matter (e.g., diffusion tensor imaging or diffusion-weighted magnetic resonance imaging [MRI]), and resting-state data (i.e., blood oxygenation level dependent [BOLD] signal collected while participants are not engaged in a task). Although these data may provide necessary information specific to an individual study, on their own they are of somewhat limited utility. Here, we make the simple yet novel proposition that these data be paired with additional behavioral data collected at a later time and outside of the context of the primary study. This approach could then support new investigations of brain structure and function that provide insight into individual differences (DeYoung, 2010; DeYoung & Gray, 2009).

In broad strokes, we propose that researchers produce a database of anatomical and resting-state scans from various task-based studies, along with the means to contact the participants who contributed these data. In this way, additional individual-difference data can be collected and paired with these neuroscience data to examine novel hypotheses. This simple idea can be implemented at a relatively small scale, such as within a single lab or across a group of researchers at a single institution. It could also, however, be implemented at a large scale across institutions, within an open-access framework. We begin by discussing the details of a small-scale implementation with the hopes that disseminating this idea will increase the frequency of personality neuroscience investigations based on large samples. We then discuss the potential for a large-scale implementation of the approach, raising a number of serious complications that will need to be overcome but also highlighting the advantages of working toward this resource.

Small-scale implementation of a personality neuroscience database

A small-scale implementation of our database idea can be undertaken by a group of researchers at a single institution or by a single lab. Assuming that the former is more likely, the first step to creating a flexible personality neuroscience database will be to recruit collaborators who are willing to contribute their neuroimaging data to the database. When recruiting potential collaborators, it is important to stress two main points: (1) that functional data linked to experimental tasks that could be reanalyzed to the detriment of the original researchers (i.e., missing out on a finding or having published results questioned) will not be contributed and (2) that by participating, they will gain access to a large data set that would be resource intensive to compile on their own, providing the opportunity to conduct new personality neuroscience studies for low additional cost. Even small groups of only two or three researchers should be able to accumulate over a hundred scans to populate the database within the first year, enough to begin doing some adequately powered correlational analyses. This estimate will vary depending on how actively each contributing lab scans participants, but if three researchers each scan 30–40 participants per year, data from approximately 200 participants will accumulate in the database within 2 years.

Once a group of researchers has committed to the idea of sharing their structural and resting-state scans, the hardware and software architecture needed to host such a database securely will need to be established. For small-scale implementations where only the contributors will be using the database, there are a number of possible solutions to this problem. Perhaps the easiest solution with respect to hardware would be to integrate the database with an already existing server system hosted by a university. Many institutions with MRI scanners already have very secure, climate-controlled, server rooms with multiple redundancies to prevent data loss. Hosting the shared data using these facilities should be simple. However, if an institution lacks such resources, at this scale even a basic Linux server in a secure room will suffice. As for the software architecture, there are again numerous solutions available. Since only a few individuals are using the database, any systematic organization of the neuroimaging data will suffice. Open-source database infrastructures for neuroimaging data already exist, such as XNAT (Marcus, Olsen, Ramaratnam, & Buckner, 2007; see also PyXNAT, Schwartz et al., 2012). These can help to ensure standardization and clear organization of data throughout the database. A data-processing pipeline such as Nipype (Gorgolewski et al., 2011) and the LONI pipeline (Dinov et al., 2009) can also be employed to standardize workflow and analysis. For the purposes of preserving anonymity, an important element of this database is that the neuroimaging data and the contact information are stored separately. Storing the contact and demographic data will require a different database infrastructure, with numerous potential solutions possible, including the open-source REDCap project (Harris et al., 2009). However, one must also be able to reassociate the neuroimaging data and contact information so that new behavioral data can be connected to the previously collected neuroimaging data. The easiest way to do this is to generate a unique identification code that will be connected to both the neuroscience data and the contact information.

Once the hardware and software infrastructure has been agreed upon, the next step is to populate the database. Since every task-based neuroimaging study also entails the collection of high-resolution anatomical scans, these are the data most likely to be contributed in the beginning. Collaborating researchers simply transfer a copy of these scans to the database during the course of their normal research. Serious consideration must be taken to preserve anonymity, with the option of removing facial information from the anatomical scans through skull-stripping, blurring, or other means (e.g., the Freesurfer defacer provided by the Biomedical Informatics Research Network [BIRN]). White matter imaging and resting-state scans are becoming a part of routine data collection during task-based studies, and these data may be contributed as well. It is possible, but not necessary, for collaborating researchers to agree to collect these data using standardized protocols, with these run on an ad hoc basis whenever there is unused time in a scanning session. In this way, participating in the database will not incur any additional cost by way of additional scan hours, but currently unused time will go toward contributing more (and perhaps more diverse forms of) data to the database. Having the protocols for these additional scans agreed upon, programmed, and ready to run will make things run more efficiently. There is currently some debate about the best techniques for collecting resting-state data that we will not attempt to resolve here, but some guidelines for collecting this kind of data have been suggested on the basis of empirical work (Van Dijk et al., 2010; cf. Handwerker, Roopchansingh, Gonzalez-Castillo, & Bandettini, 2012; Petridou, Gaudes, Dryden, Francis, & Gowland, 2013). With respect to the imaging of white matter tracts, there is again some debate, along with reasonable advice, as to how these data should best be collected and analyzed (e.g., Jones, Knösche, & Turner, 2013).

In addition to the neuroimaging data, the contact information for all participants will have to be collected in order to make it possible to acquire additional behavioral measurements that can then be paired with the neuroimaging data. The easiest way to do this is to include a question at the bottom of consent forms that asks whether participants would be willing to be contacted in the future for additional studies (for pay or course credit), followed by a line for their e-mail address. At many institutions, questions such as this are commonly included on all consent forms. However, researchers should check with their institutional review boards (IRBs) to ensure that gaining this permission for future contact is in accordance with institutional policy. It is also important for participants to actively consent to sharing their anonymized data with other researchers, including in the consent form an option for them to indicate their willingness to do so. In addition to this information, it also makes sense to present a basic demographics questionnaire for each participant to complete. At the very least, the individual's name, birthdate, and gender should be recorded. The identity information (i.e., name and birthdate) will then be stored separately from the data, but linked to it through a unique identifier. Additional basic demographic questions can always be included, such as handedness, number of siblings and birth order, and language abilities. A brief questionnaire assessing broad personality traits, such as the 44-item Big Five Inventory (John & Srivastava, 1999), would also be helpful to include, since these measures provide good coverage for a wealth of broad behavioral, affective, and cognitive tendencies (Goldberg, 1990).

After a year or so, even a small group of neuroimagers should have between 100 and 200 scanned participants who have agreed to be contacted in the future, with the neuroimaging data stored in one database, names and contact information stored in a separate database, and a system for associating the two. At this point, a novel study can be attempted with the database. As an illustration of how such a study might be conducted, consider the question of whether the volume of the medial parietal cortex, one of the brain's most highly connected processing hubs (Hagmann et al., 2008), is related to the ability to solve insight problems (Lockhart, Lamon, & Gick, 1988).[2] In order to examine this hypothesis, an online survey would be created that asks participants to solve a series of insight problems (which are characterized by sudden shifts in cognitive framing that precede awareness of the solution).[3]

An e-mail would then be sent to everyone who is in the database asking if they would be willing to complete this study for monetary remuneration. Ideally, sending this e-mail to everyone in the database will be an automated process, using a secure listserv. It is assumed that not everyone will respond and fill out the survey, but over time the database should include a large enough number of possible participants that any request will be met with a sufficient number of willing individuals. Researchers interested in collecting informant reports have found a high response rate for e-mail requests to complete short surveys, even when no remuneration is offered (e.g., above 75 %; Vazire, 2006). As well, Amazon's Mechanical Turk (mTurk) system has established that many individuals are willing to do short tasks for small sums; the current going rate for payment on mTurk is about $6 (US) per half hour. Moreover, the data that result from services such as mTurk, and online survey software (e.g., Qualtrics), appear to be of equivalent quality to those collected in-lab (Buhrmester, Kwang, & Gosling, 2011; Chuah, Drasgow, & Roberts, 2006), although additional issues, such as data security, need to be considered (Nosek, Banaji, & Greenwald, 2002). It is therefore possible that collecting additional data for 150 individuals who have already contributed high-resolution anatomical scans might cost no more than $900 (or less, if your task takes less than half an hour to complete). Paying participants who reply is possible through a number of different means, including online systems such as PayPal or e-mail money transfers. Simpler options also exist, such as having participants provide a mailing address where a check can be sent or having an office where they can pick up their payment in cash upon presentation of valid identification.

Once the data have been collected, responses to the survey regarding insight problem-solving ability will need to be associated with the corresponding anatomical scans collected previously. Some online survey services allow for unique identifiers to be embedded within the survey data (e.g., Qualtrics), which could assist with this task. Voxel-based morphometry (Ashburner & Friston, 2000) can then be employed to examine whether the proportion of gray matter in the medial parietal cortex is related to the ability to solve insight problems. This example illustrates how the kind of database we are proposing will allow valuable personality neuroscience research to be conducted using large samples (and therefore, adequate statistical power), for relatively low

[2] In this example, we consider a psychological individual difference, but the method can easily be extended to any differences between individuals, including physiological ones, such as handedness, eyeglass prescription strength, and height.

[3] An example of such a problem is as follows:

A man in a town married 20 women. He and the women are still alive, and he has had no divorces or annulments. He is not a bigamist (meaning he is not legally married to more than one woman at once), and he broke no law. How is that possible?

The solution to this problem is that the man is a priest or justice of the peace who performed these marriage ceremonies.

additional cost, while making use of neuroimaging data that have already been collected for other purposes but are typically underused. All of the personality neuroscience studies cited earlier could have been conducted for significantly less cost using this approach. Lowering the financial bar to perform these types of studies with large samples will likely have a number of positive influences on the field, including (but not limited to) conducting important replication attempts (Pashler & Harris, 2012), attenuating the inferential problems associated with low statistical power (Button et al., 2013) and maximizing the value of neuroimaging data that are relatively expensive to collect.

Note that the studies made possible with a database of this kind are not limited to online tasks. If individuals who contributed the neuroimaging data live in the area and are willing to come visit a lab, they can easily complete tasks and measurements in person. This might be necessary for studies that require other forms of physiological measurement (e.g., heart rate, cortisol levels, skin conductance, gait analyses, more complex measurement of perception or sensation) or behavior (e.g., social interaction tasks). Studies also need not be limited to a single time-point. With such a database, it will be entirely possible to perform longitudinal studies or randomized control intervention studies that also incorporate an investigation of neuroanatomy and functional connectivity. Moreover, it is not uncommon for an individual to participate in more than one fMRI study at a particular institution. Keeping track of these individuals will allow you to examine possible changes in brain structure and function over time and how these changes might relate to other experiences (e.g., courses taken at a university, drug use, media use) or individual differences (e.g., trait personality). Thus, a growing database will need to take into account individuals who have been scanned previously. Asking whether a person has been scanned previously in the demographics questionnaire should be sufficient to alert researchers that this individual might already have an entry in the database.

Even with respect to a small-scale implementation of this idea, with only a few researchers contributing to and employing the database, additional issues are likely to arise as the database continues to grow. First, it will be important to ensure that potential participants are not bombarded with too many requests for studies. A possible solution is to include an algorithm that ensures that participants are not contacted more than 3 or 4 times a year. A system could also be established in which participants indicate the frequency with which they are comfortable being contacted regarding possible study participation. Listserv or newsletter services such as JISCMail allow recipients to opt out of e-mails or choose to receive an aggregate of all messages sent each week or each month (a digest version). It might also be possible to establish a Web site where each person who has contributed neuroimaging data can browse available studies whenever he or she wishes. These ideas are not mutually exclusive, and one can easily imagine a combination of any number of these possible solutions being successful. At this scale of implementation, solutions to most problems can be decided upon by the collaborating researchers who may, for example, agree that each contributing lab pursue only one or two such studies per year in order to moderate usage.

Large-scale implementation of a personality neuroscience database

At a small-scale, this database idea is a simple way for small groups of researchers to produce a flexible resource that permits adequately-powered personality neuroscience investigations. The main issues of constructing, maintaining, and using the database at this scale can be solved in a myriad of different ways, and it would not be surprising if entirely novel solutions separate from those mentioned here are employed. However, it is possible for this database idea to be scaled up to provide an even greater resource accessible to a large number of researchers. In this section, we outline what a large-scale implementation of a flexible personality neuroscience database might look like.

In order to implement this database idea at a large scale, one possibility is to allow contributions from many different researchers, including those at other institutions. Doing so would allow the database to increase in size, but allowing multisite participation brings a number of other hurdles that must be overcome. Although, anatomical scans do not differ widely in format, these data and the associated contact information must be transferred securely and with the agreement of the IRBs for all participating institutions. Moreover, a much larger and more secure hosting infrastructure would be required, with a dedicated server room being the ideal solution.

Key to a large-scale implementation of this idea will be the issue of access. In order to maximize utility, any qualified researcher should be able to access and employ the database, as is the case for other major data release initiatives (e.g., ADNI, Mueller et al., 2005; OASIS, Marcus, Wang, Parker, Csernansky, Morris, & Buckner 2007; BIRN, Zou et al., 2005; HCP, Marcus et al., 2011). Enabling this access raises a number of important additional issues that are not present in the small-scale implementation. Although the number of participants can grow if the database is opened up to multiple contributors, the number of study requests made to participants must still be moderated so that they do not create a burden on those who have contributed the individual data. Before opening the database to all researchers, then, it might make sense to limit access in some way, perhaps to those who also contributed data.

In addition to the "who" question of access, one must consider the question of "how." Because very large amounts of data will exist and the database will be constantly expanding, it is not likely to be feasible to allow the database to be downloaded or transferred between locations. A more parsimonious solution is to allow users from around the world to view and analyze the data while it remains in a single location, using an Internet portal. Keeping the data at a single site allows for greater control and management of the data. This also accords simplicity, in that only one IRB will be involved in overseeing the ethical issues pertaining to maintenance and access of the database; the same IRB will also oversee the conduct of the associated studies. In order to host the data at a single site but still permit wide access, an analysis portal is required. Such a portal would allow for analysis procedures to be initiated and applied to the combined behavioral and neuroimaging data remotely, without the need for the data to be copied or transferred to other locations. One way to achieve this might be to employ the LONI pipeline and its Web Start functionality (Dinov et al., 2009). In this way, many people may share access to the data, while keeping the data anonymized and secure. As an additional advantage, a system of this sort might also allow for a complete log of the analyses to be recorded. Making this log public, perhaps in the form of a URL provided with each article resulting from the database, will result in greater transparency and replicability (addressing two additional issues for research).

Once such a system has been implemented alongside a flexible personality neuroscience database, we believe that maintaining and operating such a system will require the assistance of full-time staff such as administrative support and analysis support. Salaries for these staff members could be paid by various sources, but the most likely is a federal research grant. These staff members would manage contact with interested researchers, helping them to pose their surveys and tasks to participants who have contributed neuroimaging data to the database and assist with the analysis process. In Table 1, we outline a possible workflow for a large-scale implementation of our database idea. To begin with, the personality neuroscience database is populated by the organizers (step 1). Users then contribute neuroimaging data (anatomical and/or resting state) that are paired with contact information for those participants (step 2). These data are screened for quality and proper anonymization of participants, then added to the database. Users wishing to conduct a study then submit a proposal that is evaluated by the database organizers and staff for feasibility and appropriateness (step 3). At this stage, organizers can ensure that proposals are not redundant with other projects already completed or underway. Studies selected to move forward would be proposed to the IRB of the institution hosting the database. These ethics proposals would be prepared by the users with assistance from the staff, but the actual submission would be done internally, with the database owners acting as principal investigators and the users as collaborators. Although this arrangement is necessary to allow a local IRB to oversee the research process, we are recommending as a form of best practice that only the users be listed as authors on resultant publications, with the database organizers and funding sources acknowledged in the acknowledgments. Once the study has been approved by the IRB, the users would program their online study and submit a link to the database staff. These staff would then make the study link available to database participants (e.g., posted to a listserv or a Web site) (step 4). Although users provide the funds for participant reimbursement, coordination of this payment is provided by the database organizers to maintain participant anonymity. Once data collection is complete, on the basis of a previously agreed upon timeline or target sample size, these new data will be entered into the database and associated with the contributing participant by the staff. Users will then have the opportunity to begin analyzing these data, with exclusive access to this additional data for a set period such as 2 years (step 5). After this time, the additional data will be made available to all users of the database. Using an Internet portal, users will examine how the newly collected data relate to the previously collected neuroimaging data, perhaps in conjunction with the additional behavioral data in the database. A log of all analysis steps will be recorded, and a URL generated to provide public access to this log. Users will be encouraged to publish this URL along with their data, in the spirit of transparency and support of replicability. Ideally, no fee or other compensation would be required for database access, although in the absence of external funding or support, some fee may be required to help support the salary of the staff associated with maintaining the database. After the 2-year grace period expires and the newly collected behavioral data become publicly available, we recommend that the users who originally proposed the study should have no claim to ownership of these data. Other users who employ these data in their own analyses should not be obligated to include anyone else as authors on their papers—not the database organizers nor the other users who contributed data that are incorporated into a subsequent analysis. In our opinion, an open data-sharing approach that is based on these principles would best serve the field. There is, however, serious debate on this complicated issue, and individual database organizers should consider all sides before adopting a particular policy regarding authorship (Hurko et al., 2012; Rohlfing & Poline, 2012).

Once the database grows large enough, access can be opened up to include all qualified researchers, including those who did not contribute neuroimaging data. In a future neuroscience where one or more large-scale implementations of a flexible personality database exist, a professor at a small liberal arts college can teach his or her students about brain–behavior relations by having them propose and then test a hypothesis in a large data set, for very little cost. Personality neuroscience studies involving hundreds of participants will become the norm, alleviating the need to wonder how low

Table 1 Workflow for a large-scale implementation of a flexible personality neuroscience database

Step	Organizer/Staff	Personality Neuroscience Database	Contact Database	Participant	User
1	Conducts task-based studies			Participates in research	
		Neuroimaging data added	Contact details added		
2					Contributes neuroscience data with contact details
		Neuroimaging data added	Contact details added		
3					Proposes study
	Evaluates proposal and obtains ethical approval				Programs online study
4	Contacts participants for online study				
	Coordinates compensation of participants	Behavioral data added, associated with neuroscience data		Participates in research	
5	Assists with analysis as needed				Analyzes neuroscience and behavioral data.
		After 2 years newly collected behavioral data becomes available to other users			Publishes results with link to analysis log

statistical power impacts interpretation of the results. The end result of a major, open-access database like what we are proposing would be that large-scale analyses of how individual differences in behavior and cognition relate to brain structure and function will become relatively inexpensive, increasing the number and diversity of such studies. What we propose will hopefully lead to a democratization of neuroscience, with financial barriers reduced to that of a standard behavioral study and neuroscience methods extended to scientists without geographic or financial access to an MRI scanner. Doing so will allow for major steps to be taken in neuroscience, bringing us much closer to understanding how brain and behavior relate.

Problems with our solution for the problem of power in personality neuroscience

There are a number of obstacles to the creation of a flexible personality neuroscience database of the form that we are proposing. For one, some may wonder whether this proposal is made redundant by the existence of other publicly available data sets that also combine neuroimaging data with personality and behavioral data, such as IMAGEN (Thyreau et al., 2012), the Human Connectome Project (Marcus et al., 2011), the 1000 Functional Connectomes Project (Biswal et al., 2010), and the Alzheimer's Disease Neuroimaging Initiative (ADNI; Wyman et al., 2013). Each of these important data-sharing initiatives provides qualified researchers with access to hundreds of participants who have both contributed neuroimaging data and completed other measures, including measures of trait personality. However, none of these databases allow for additional data to be collected from the participants. In other words, the questions that can be investigated with these data sets are limited by the measures included in the current data set. The flexible database framework that we are proposing, however, allows for additional data to be collected from the participants who contributed the neuroimaging data, which permits the investigation of new questions not originally conceived when the participants were first scanned. This prospective flexibility increases the utility of our proposed database, since it allows for novel questions to be asked and ensures that the utility of the original neuroimaging data is not exhausted after repeated reanalysis. Moreover, because this personality database idea can be usefully implemented at a small scale (i.e., small groups of neuroimagers or a single lab) using inexpensive resources, multiple such databases can be created and can coexist. For small-scale implementations, our proposal relies upon a very simple idea: aggregate structural and resting-state data from current neuroimaging studies, gain permission to contact past participants in the future, then collect more data from these participants to test novel hypotheses when combined with the neuroimaging data. Although simple, if this idea is adopted by researchers, it will make personality neuroscience investigations significantly less expensive to conduct, which in turn will hopefully increase the number and quality of these types of studies. At a small scale, this idea is much easier to instantiate than the major data-sharing initiatives that have been previously attempted, since access issues are limited to the small group of collaborating researchers. Admittedly, a large-scale implementation of this idea would be far more difficult to produce, although it brings a host of separate benefits to the field. That said, this framework brings a number of important benefits to the field even when only implemented at a small scale.

Another question that is often asked whenever a database project is proposed is why one believes this idea will be possible in light of past experiences with databases that have encountered difficulties (e.g., fMRIDC, Van Horn & Gazzaniga, 2012; Neurogenerator, Roland et al., 2001). For one, although these past databases did encounter difficulties, they were successful in supporting a number of new publications (Van Horn & Gazzaniga, 2012). Second, there have been numerous other database ideas that continue to support new investigations (e.g., the ADNI [Mueller et al., 2005], the 1000 Functional Connectomes Project, and the International Neuroimaging Datasharing Initiative [fcon_1000.projects.nitrc.org]), including the new Human Connectome Project that intends to publicly release petabytes of data (Marcus et al., 2011). So the struggles experienced by some past data-sharing efforts cannot alone be a reason not to pursue new data-sharing ideas, especially in light of the growing success of new efforts. Lastly, our database idea differs from past initiatives in that it can be implemented at a small scale among groups of researchers where the hurdles to contributing and sharing are minor. More important, even at a large scale, our proposal involves ease for contributors and is relatively nonthreatening. What is being requested involves minimal additional effort and little additional cost on the part of the contributing researchers. High-resolution anatomical scans are always collected as part of any neuroimaging study, and the proposed database will not require any other information beyond the participant's contact information and minimal demographic information, including sex and age. Specifically, there is no need to describe the task or any other details of the study, since the database will not require task-based functional (i.e., BOLD) data, nor will the contributed data need to be altered to suit a particular format before submission. By not asking for task-related functional MRI data, there is also little risk that someone will reanalyze the data in a way that might call the original researchers' published results into question, which is a concern that can prevent researchers from sharing their data. Moreover, the data requested are unlikely to be reanalyzed by the contributing researcher to test novel hypotheses, which means that there is little chance of a researcher missing out on an opportunity to report a new finding. The analytic approach supported by this database requires large sample sizes, larger than what most single researchers will be able to muster in a short time span. Contributing data to a flexible personality neuroscience database will entail little additional work beyond what is already being done for individual neuroimaging studies, and contribution will entail no foreseeable risks for the researcher. The reward of being able to conduct novel personality neuroscience studies with very large sample sizes for low cost, however, is huge.

Another issue is whether one can properly study brain–behavior associations when the brain data are collected at a different time from the behavioral data. This is a valid and intriguing point. Given evidence that brain structure can change in response to age (Good et al., 2001) and experience (Draganski et al., 2004; Woollett & Maguire, 2011), these are legitimate concerns. What is important to point out, however, is that these changes are unlikely to be confounded with key variables of interest and will, therefore, simply introduce noise into the analysis. In other words, if any associations are found between brain structure and an individual-difference variable, they must be strong enough to be robust to any noise introduced by nonconcurrent assessment. The variability in brain structure introduced by a time delay between imaging the neuroanatomy and measuring some behavioral trait will not be systematic or form a confound with key variables of interest, so it will not produce false positives. Additionally, basic traits are relatively stable for adults (Roberts, Walton, & Viechtbauer, 2006; Terracciano, McCrae, & Costa, 2010). For example, the Eugene-Springfield Community Sample (Goldberg, 1999) has demonstrated that even when measures are collected from the same individuals over a span of time as long as 10 years, reliable associations can be uncovered that replicate in samples where all measures were collected concurrently (e.g., DeYoung, Quilty, & Peterson, 2007). Moreover, duration of time between scanning and behavioral data collection can always be used as a covariate in analyses to take into account the time between measurements. One real danger with delays between neuroimaging and behavioral assessment is that real effects may be attenuated by whatever noise is introduced into the analysis by the delay. Nonetheless, given the fact that creating this database and pursuing novel questions will not be costly, coupled with the advantages of operating with sufficient statistical power to detect even small effects, the benefits would seem to outweigh the costs.

An additional problem to consider is attrition. To what degree will researchers be able to remain in contact with the participants who contributed the original neuroimaging data? One concern might be that undergraduate participants will be unreachable after graduation, which is likely if these participants provide University e-mail accounts that are closed following graduation. However, as e-mail becomes a primary mode of contact and communication, many universities are now offering lifelong e-mail addresses to their students. For universities who fail to do so, most students also maintain an e-mail address that they intend to keep for life, since changing accounts involves substantial hassle. To minimize the problem of students' providing e-mail addresses that they will stop using upon graduation, those willing to be contacted in the future can be explicitly encouraged to provide an e-mail address they intend to keep for life. Moreover, the long history of longitudinal research in psychology and other disciplines has already established that although attrition is a problem, it is not an insurmountable one.

Another concern is whether the flexible and extensible nature of the proposed database means that anonymity will become harder to maintain as more and more data are collected. This is a genuine concern, in that a large amount of demographic and geographic data could hypothetically be employed to infer identity. However, in practice, the variables of interest to researchers are unlikely to provide enough information to support accurate inferences regarding identity (e.g., task performance or self-reports of traits and motivations). That said, this is an important issue, and care should be taken to ensure that the additional data collected are unlikely to provide enough information to deanonymize participants. This can be taken into consideration when study proposals are evaluated. In conjunction, data-use agreements should explicitly forbid attempting to identify any research participant.

Conclusions

In proposing this idea for how a flexible and extensible personality neuroscience database might be created, we hope that researchers will consider implementing this idea on a small scale. Doing so will require little by way of additional resources and cost but will permit novel investigations of how brain and behavior relate based on large samples affording adequate statistical power. Although a large-scale implementation of this idea would entail a number of additional difficulties that will require careful solutions, it would also bring with it additional benefits, such as greater access and democratization of neuroscience research. It is our hope that these ideas promote greater data-sharing among researchers, since this is the simplest way that issues of low statistical power can be overcome for personality neuroscience investigations.

Author Notes RAM gratefully acknowledges the support of York University's Dean's Health Research Catalyst Award, which made producing this article possible. Dr. Jolynn Pek and Kevin Desimone are also thanked for their assistance with the manuscript.

References

Adelstein, J. S., Shehzad, Z., Mennes, M., DeYoung, C. G., Zuo, X.-N., Kelly, C., ... Milham, M. P. (2011). Personality is reflected in the brain's intrinsic functional architecture. *PLoS ONE, 6,* e27633.

Ahadi, S., & Diener, E. (1989). Multiple determinants and effect size. *Journal of Personality and Social Psychology, 56,* 398–406.

Andrews-Hanna, J. R., Reidler, J. S., Huang, C., & Buckner, R. L. (2010). Evidence for the default network's role in spontaneous cognition. *Journal of Neurophysiology, 104,* 322–335.

Ashburner, J., & Friston, K. J. (2000). Voxel-based morphometry—the methods. *NeuroImage, 11,* 805–821.

Banissy, M. J., Kanai, R., Walsh, V., & Rees, G. (2012). Inter-individual differences in empathy are reflected in human brain structure. *NeuroImage, 62,* 2034–2039.

Bickart, K. C., Wright, C. I., Dautoff, R. J., Dickerson, B. C., & Barrett, L. F. (2011). Amygdala volume and social network size in humans. *Nature Neuroscience, 14,* 163–164.

Biswal, B. B., Mennes, M., Zuo, X. N., Gohel, S., Kelly, C., Smith, S. M., ... Milham, M. P. (2010). Toward discovery science of human brain function. *Proceedings of the National Academy of Sciences, 107,* 4734–4739.

Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science, 6,* 3–5.

Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, S. J., & Munafo, M. R. (2013). Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience, 14,* 1–14.

Chuah, S. C., Drasgow, F., & Roberts, B. W. (2006). Personality assessment: Does the medium matter? No. *Journal of Research in Personality, 40,* 359–376.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.

DeCarli, C., Murphy, D. G. M., Tranh, M., Grady, C. L., Haxby, J. V., Gillette, J. A., ... Schapiro, M. B. (1995). The effect of white-matter hyperintensity volume on brain structure, cognitive performance, and cerebral metabolism in healthy adults. *Neurology, 45,* 2077–2084.

DeYoung, C. G. (2010). Personality neuroscience and the biology of traits. *Social and Personality Psychology Compass, 4,* 1165–1180.

DeYoung, C. G., & Gray, J. R. (2009). Personality neuroscience: Explaining individual differences in affect, behavior, and cognition. In P. J. Corr & G. Matthews (Eds.), *Cambridge handbook of personality* (pp. 323–346). New York: Cambridge University Press.

DeYoung, C. G., Hirsh, J. B., Shane, M. S., Papademetris, X., Rajeevan, N., & Gray, J. R. (2010). Testing predictions from personality neuroscience: Brain structure and the Big Five. *Psychological Science, 21,* 820–828.

DeYoung, C. G., Quilty, L. C., & Peterson, J. B. (2007). Between facets and domains: 10 Aspects of the Big Five. *Journal of Personality and Social Psychology, 93,* 880–896.

Dinov, I., Van Horn, J. D., Lozev, K. M., Magsipoc, R., Petrosyan, P., Liu, Z., ... Toga, A. W. (2009). Efficient, distributed and interactive neuroimaging data analysis using the LONI pipeline. *Frontiers in Neuroinformatics, 3,* 22.

Doucet, G., Naveau, M, Petit, L., Zago, L., Crivello, F., Jobard, G., ... Joliot, M. (2012). Patterns of hemodynamic low-frequency oscillations in the brain are modulated by the nature of free thought during rest. *NeuroImage, 59,* 3194–3200.

Draganski, B., Gaser, C., Busch, V., Schuierer, G., Bogdahn, U., & May, A. (2004). Neuroplasticity: Changes in grey matter induced by training. *Nature, 427,* 311–312.

Efron, B., & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science, 1,* 54–75.

Efron, B., & Tibshirani, R. (1993). *An introduction to the bootstrap*. New York: Chapman & Hall.

Fraley, R. C., & Marks, M. J. (2007). The null hypothesis significance testing debate and its implications for personality research. In R. W. Robins, R. C. Fraley, & R. Krueger (Eds.), *Handbook of research methods in personality psychology* (pp. 149–169). New York: Guilford Press.

Goldberg, L. R. (1990). An alternative "description of personality": The Big-Five factor structure. *Journal of Personality and Social Psychology, 59,* 1216–1229.

Goldberg, L. R. (1999). A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models. In I. Mervielde, I. Deary, F. De Fruyt, & F. Ostendorf (Eds.), *Personality psychology in Europe, Vol. 7* (pp. 7–28). Tilburg, the Netherlands: Tilburg University Press.

Good, C. D., Johnsrude, I. S., Ashburner, J., Henson, R. N. A., Friston, K. J., & Frackowiak, R. S. J. (2001). A voxel-based morphometric study of ageing in 465 normal adult human brains. *NeuroImage, 14,* 21–36.

Gonzalez-Castillo, J., Saad, Z., Handwerker, D., Inati, S., Brenowitz, N., and Bandettini, P. (2012). Whole-brain, time-locked activation with simple tasks revealed using massive averaging and model-free analysis. *Proceedings of the National Academy of Sciences, 109*(14), 5487–5492. doi:10.1073/pnas.1121049109

Gorgolewski, K., Burns, C. D., Madison, C., Clark, D., Halchenko, Y. O., Waskom, M. L., & Ghosh, S. S. (2011). Nipype: A flexible, lightweight and extensible neuroimaging data processing framework in Python. *Frontiers in Neuroinformatics., 5,* 13.

Green, A. E., Munafo, M. R., DeYoung, C. G., Fossella, J. A., Fan, J., & Gray, J. R. (2008). Using genetic data in cognitive neuroscience: From growing pains to genuine insights. *Nature Reviews Neuroscience, 9,* 710–720.

Gur, R. C., Turetsky, B. I., Matsui, M., Yan, M., Bilker, W., Hughett, P., & Gur, R. E. (1999). Sex differences in brain gray and white matter in healthy young adults: Correlations with cognitive performance. *Journal of Neuroscience, 19,* 4065–4072.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. M., & Spore, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS Biology, 6*, e159.

Handwerker, D. A., Roopchansingh, V., Gonzalez-Castillo, J., & Bandettini, P. A. (2012). Periodic changes in fMRI connectivity. *NeuroImage, 63,* 1712–1719.

Harris, P. A., Taylor, R., Thielke, R., Payne, J., Gonzalez, N., & Conde, J. G. (2009). Research electronic data capture (REDCap): A metadata-driven methodology and workflow process for providing translational research informatics support. *Journal of Biomedical Informatics, 42,* 377–381.

Hemphill, J. F. (2003). Interpreting the magnitude of correlation coefficients. *American Psychologist, 58,* 78–79.

Holmes, A. J., Lee, P. H., Hollinshead, M., Bakst, L., Roffman, J. L., Smoller, J. W., & Buckner, R. L. (2012). Individual differences in amygdala-medial prefrontal anatomy link negative affect, impaired social functioning, and polygenic depression risk. *Journal of Neuroscience, 32,* 18087–18100.

Hu, X., Erb, M., Ackermann, H., Martin, J. A., Grodd, W., & Reiterer, S. M. (2011). Voxel-based morphometry studies of personality: Issue of statistical model specification–effect of nuisance covariates. *NeuroImage, 54,* 1994–2005.

Hurko, O., Black, S. E., Doody, R., Doraiswamy, P. M., Gamst, A., Kaye, J., ... the ADNI Data and Publication Committee. (2012). The ADNI Publication Policy: Commensurate recognition of critical contributors who are not authors. *NeuroImage, 59,* 4196–4200.

Jennings, R. G., & Van Horn, J. D. (2012). Publication bias in neuroimaging research: Implications for meta-analyses. *Neuroinformatics, 10,* 67–80.

John, O. P., & Srivastava, S. (1999). The Big Five Trait taxonomy: History, measurement, and theoretical perspectives. In L. A. Pervin & O. P. John (Eds.), *Handbook of personality: Theory and research* (pp. 102–138). New York: Guilford Press.

Jones, D. K., Knösche, T. R., & Turner, R. (2013). White matter integrity, fiber count, and other fallacies: The do's and don'ts of diffusion MRI. *NeuroImage, 73,* 239–254.

Kanai, R., Bahrami, B., & Rees, G. (2010). Human parietal cortex structure predicts individual differences in perceptual rivalry. *Current Biology, 20,* 1626–1630.

Kanai, R., Bahrami, B., Roylance, R., & Rees, G. (2012). Online social network size is reflected in human brain structure. *Proceedings of the Royal Society B: Biological Sciences, 279,* 1327–1334.

Kapogiannis, D., Sutin, A., Davatzikos, C., Costa Jr., P., & Resnick, S. (2012). The five factors of personality and regional cortical variability in the Baltimore longitudinal study of aging. *Human Brain Mapping*. doi:10.1002/hbm.22108

Kilpatrick, L., Zald, D. H., Pardo, J. V., & Cahill, L. (2006). Sex-related differences in amygdala functional connectivity during resting conditions. *NeuroImage, 30,* 452–461.

Koyama, M. S., Di Martino, A., Zuo, X., Kelly, C., Mennes, M., Jutagir, D. R., ... Milham, M. P. (2011). Resting-state functional connectivity indexes reading competence in children and adults. *Journal of Neuroscience, 31,* 8617–8624.

Lee, W. C., & Rodgers, J. L. (1998). Bootstrapping correlation coefficients using univariate and bivariate sampling. *Psychological Methods, 3,* 91–103.

Lockhart, R. S., Lamon, M., & Gick, M. (1988). Conceptual transfer in simple insight problems. *Memory & Cognition, 16,* 36–44.

Marcus, D. S., Harwell, J., Olsen, T., Hodge, M., Glasser, M. F., Prior, F., ... Van Essen, D. C. (2011). Informatics and data mining tools and strategies for the human connectome project. *Frontiers in Neuroinformatics, 5,* 4.

Marcus, D. S., Olsen, T., Ramaratnam, M., & Buckner, R. L. (2007). The Extensible Neuroimaging Archive Toolkit (XNAT): An informatics platform for managing, exploring, and sharing neuroimaging data. *Neuroinformatics, 5,* 11–34.

Marcus, D. S., Wang, T. H., Parker, J., Csernansky, J. G., Morris, J. C., & Buckner, R. L. (2007). Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *Journal of Cognitive Neuroscience, 19,* 1498–1507.

Meyer, G. J., Finn, S. E., Eyde, L., Kay, G. G., Moreland, K. L., Dies, R. R., ... Reed, G. M. (2001). Psychological testing and psychological assessment: A review of evidence and issues. *American Psychologist, 56,* 128–165.

Mills, K. L., Lalonde, F., Clasen, L., Giedd, J. N., & Blakemore, S. -J. (2012) Developmental changes in the structure of the social brain in late childhood and adolescence. *Social Cognitive and Affective Neuroscience*. doi:10.1093/scan/nss113

Mueller, S. G., Weiner, M. W., Thal, L. J., Petersen, R. C., Jack, C. R., Jagust, W., ... Becket, L. (2005). Ways toward an early diagnosis in Alzheimer's disease: The Alzheimer's Disease Neuroimaging Initiative (ADNI). *Alzheimer's Dementia, 1,* 55–66.

Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). E-Research: Ethics, security, design, and control in psychological research on the Internet. *Journal of Social Issues, 58,* 161–176.

Parkinson, C. & Wheatley, T. (2012). Relating anatomical and social connectivity: White matter microstructure predicts emotional empathy. *Cerebral Cortex*. doi:10.1093/cercor/bhs347

Pashler, H., & Harris, C. R. (2012). Is the replicability crisis overblown? Three arguments examined. *Perspectives in Psychological Science, 7,* 531–536.

Petridou, N., Gaudes, C. C., Dryden, I. L., Francis, S. T., & Gowland, P. A. (2013). Periods of rest in fMRI contain individual spontaneous events which are related to slowly fluctuating spontaneous activity. *Human Brain Mapping, 34,* 1319–1329.

Poldrack, R. A. (2012). The future of fMRI in cognitive neuroscience. *NeuroImage, 62,* 1216–1220.

Roberts, B. W., Walton, K.E. & Viechtbauer, W. (2006). Patterns of mean-level change in personality traits across the life course: A meta-analysis of longitudinal studies. *Psychological Bulletin, 132,* 3–27.

Rohlfing, J., & Poline, J.-B. (2012). Why shared data should not be acknowledged on the author byline. *NeuroImage, 59,* 4189–4195.

Roland, P., Svensson, G., Lindeberg, T., Risch, T., Baumann, P., Dehmel, A., ... Zilles, K. (2001). A database generator for human brain imaging. *Trends in Neurosciences, 24,* 562–564.

Schwartz, Y., Barbot, A., Thyreau, B., Frouin, V., Varoquaux, G., Siram, A., ... Poline, J. B. (2012). PyXNAT: XNAT in Python. *Frontiers in Neuroinformatics, 6,* 12.

Strube, M. J. (1991). Multiple determinants and effect size: A more general method of discourse. *Journal of Personality and Social Psychology, 61,* 1024–1027.

Terracciano, A., McCrae, R. R., & Costa, P. T., Jr. (2010). Intra-individual change in personality stability and age. *Journal of Research in Personality, 44,* 31–37.

Thyreau, B., Schwartz, Y., Thirion, B., Frouin, V., Loth, E., Vollstädt-Klein, S., ... Poline, J. (2012). Very large fMRI study using the IMAGEN database: Sensitivity–specificity and population effect modeling in relation to the underlying anatomy. *NeuroImage, 61,* 295–303. doi:10.1016/j.neuroimage.2012.02.083

Van Dijk, K. R. A., Hedden, T., Venkataraman, A., Evans, K. C., Lazar, S. W., & Buckner, R. L. (2010). Intrinsic functional connectivity as a tool for human connectomics: Theory, properties, and optimization. *Journal of Neurophysiology, 103,* 297–321.

Van Horn, J. D. & Gazzaniga, M. S. (2012). Why share data? Lessons learned from the fMRIDC. *NeuroImage*. doi:10.1016/j.neuroimage.2012.11.010

Vazire, S. (2006). Informant reports: A cheap, fast, and easy method for personality assessment. *Journal of Research in Personality, 40,* 472–481.

Woollett, K., & Maguire, E. A. (2011). Acquiring “the Knowledge” of London’s layout drives structural brain changes. *Current Biology, 21,* 2109–2114.

Wyman, B. T., Harvey, D. J., Crawford, K., Bernstein, M. A., Carmichael, O., Cole, P. E., et al. (2013). Standardization of analysis sets for reporting results from ADNI MRI data. *Alzheimer's & Dementia, 9,* 332–337.

Yarkoni, T. (2009). Big correlations in little studies: Inflated fMRI correlations reflect low statistical power. Commentary on Vul, et al. (2009). *Perspectives on Psychological Science, 4,* 294–298.

Yarkoni, T., & Braver, T. S. (2010). Functional neuroimaging approaches to individual differences in working memory and executive control: Conceptual and methodological considerations. In Gruszka, Matthews, & Szymura (Eds.), *Handbook of individual differences in cognition* (pp. 87–107). Berlin: Springer.

Yarkoni, T., Poldrack, R. A., Van Essen, D. C., & Wager, T. D. (2010). Cognitive neuroscience 2.0: Building a cumulative science of human brain function. *Trends in Cognitive Sciences, 14,* 489–496.

Yu, R. (2012). Regional white matter volumes correlate with delay discounting. *PLoS One, 7*(2), e32595.

Zou, K. H., Greve, D. N., Wang, M., Pieper, S. D., Warfield, S. K., White, N. S., ... for the FIRST BIRN Research Group. (2005). Reproducibility of functional MR imaging: Preliminary results of prospective multi-institutional study performed by Biomedical Informatics Research Network. *Radiology, 237,* 781–789.

Zuckerman, M. (2005). *Psychobiology of personality* (Revised edth ed.). New York, NY: Cambridge University Press.