

Personality Psychology

An open access textbook

Patrick Durkee

August 2026

Table of contents

Welcome	5
Contents	5
License	5
1 Foundational Frameworks	6
Navigating Questions and Explanations	6
Connecting to Personality	12
Wrapping Up and Building on These Foundations	15
References	16
2 Past Perspectives	17
Is Personality Ancient?	18
A Nonexhaustive History of Historical Personality Perspectives	19
Conclusion	24
References	24
3 Contemporary Perspectives	25
Too Many Theories, Not Enough Data	25
The Dictionary as Data	25
People as Data	26
The Significance of the Big Five	30
A Sixth Personality Factor?	30
So How Many Personality Dimensions Are There?	31
Personality Models Are Useful Fictions	31
References	32
4 Measurement Matters	33
The Relationship Between Reliability and Validity	33
Reliability Assessment	34
Validity Assessments	38
What's the Verdict? Do Personality Inventories Measure What They Intend To?	43
References	44
5 Personality Power	45
The Person-Situation Debate	45
Walter Mischel's Critique of Personality	46

The Power of the Situation	47
The Power of Personality	48
Resolving the Person Situation Debate	49
References	49
6 Evolution and Personality	51
Darwin's Big Ideas	51
Products of Evolution	54
The Power of Evolutionary Theory	55
Applying Darwin's Ideas to Humans: Evolutionary Psychology	56
7 Genetics and Personality	58
Notes of Caution	58
Nature and Nurture(s)	59
Estimating the Contributions of Genetics and Environments with Twin Studies	60
Getting at the Genes Through Molecular Genetics	64
Conclusions About Genetics and Personality Heritability	70
Potential Evolutionary Explanations	70
References	72
8 Physiology and Personality	74
Brains	74
Bodies	81
Conclusions About Physiology and Personality	86
References	86
9 Emotions and Personality	89
An Evolutionary Framework for Studying Emotions	89
Individual Differences in Emotions	93
Reverse Engineering Personality Constructs via Emotions	94
Emotions Provide Insight into Lexical Personality Constructs	96
References	96
10 Motivation and Personality	98
What Are Motivations?	98
What Goals Are Motivations For?	100
Closing Thoughts About Motivations and Personality	106
References	106
11 Development and Personality	107
Temperament versus Personality	107
The Little Six?	107
General Principles of Personality Development	108
Normative Trends Aren't "One Size Fits All"	111

References	111
12 Culture and Personality	112
Common Pitfalls to Avoid	113
Elements of Culture	114
Dimensions of Culture	115
Interplay Between Culture and Personality	117
Ecological Factors	120
Cultural Variation in Personality Structure	121
Closing Thoughts About Culture and Personality	123
References	124
13 Personality Change	126
Many People Want to Change Their Personality	126
From Personality States to Personality Traits	127
Participant Attributes That May Affect Change	128
Intervention Attributes That May Affect Change	129
Evidence from Psychotherapy and Clinical Intervention	130
Evidence from Volitional Personality Change Studies	131
How Much Change Is Possible?	131
Closing Thoughts About Personality Change	132
References	133

Welcome

[Start reading](#) →

Contents

1. [Foundational Frameworks](#)
2. [Past Perspectives](#)
3. [Contemporary Perspectives](#)
4. [Measurement Matters](#)
5. [Personality Power](#)
6. [Evolution and Personality](#)
7. [Genetics and Personality](#)
8. [Physiology and Personality](#)
9. [Emotions and Personality](#)
10. [Motivation and Personality](#)
11. [Development and Personality](#)
12. [Culture and Personality](#)
13. [Personality Change](#)

License

Except where otherwise noted, this work is licensed under a [Creative Commons Attribution-NonCommercial 4.0 International License](#).

You are free to read, share, and adapt this book for noncommercial purposes, as long as you give credit and indicate any changes. Commercial rights are reserved.

Corrections and suggestions are welcome by [email](#).

1 Foundational Frameworks

Studying psychology can feel a bit disorienting. Many undergraduate courses and research topics in psychology appear to be only loosely connected. The psychological literature is filled with myriad empirical findings, many of which are not even be reliably repeatable (Anvari & Lakens, 2018). Professors and researchers defend “theories” that feel more like confidently stated hunches about psychological phenomena than well-developed explanations and descriptions (Eronen & Bringmann, 2021).

Students are often left to try to make sense of this mishmash without an overarching theoretical framework to guide us in interpreting and parsing the information we are given. It’s like trying to put together a puzzle without knowing what it should look like at the end. In this chapter, we’ll discover theoretical frameworks that can provide us some hints at how we might go about putting the puzzle together — or at least figure out what jigsaw pieces we are looking for.

Navigating Questions and Explanations

To make progress in science, we need to be able to ask good research questions and develop good explanations of the phenomena we are interested in. But there are many ways we could ask questions and develop explanations. At a very broad level, we can ask *why a thing exists* and we can ask *how a thing works*. For example, if we want to develop a scientific understanding of vision, we can ask why eyes exist and we can ask how eyes work with our mind-brain to produce the experience of sight. The first question is often called an *ultimate* question. The second question is an example of a *proximate* question. Within each of these broad questions, we can also ask many more specific questions: How does the lens refract light into the retina? How is light converted into electrical signals that can be interpreted by our occipital lobes? What does vision accomplish? And so on.

Although we’ll see that these types of questions provide somewhat different pieces of an explanation we want to assemble, they are often mistaken by students (and researchers) as *competing* explanations. This can lead to confusion—and even hostility—when people conflate discussion of one particular question or explanation with claiming that others are unimportant or irrelevant. But these questions and explanations are actually *complimentary*. Rather than one type of question or explanation diminishing the importance of another, all types of explanations are necessary to fully understand the phenomenon. We might have our favorite type, but that doesn’t mean others are unimportant.

This ultimate versus proximate distinction can be broken down further to include questions (and explanations) about *phylogeny*, *function*, *ontogeny*, and *mechanism*, which are often referred to as **Tinbergen’s Four Questions** after Nicolaas Tinbergen, a prominent biologist who wrote about them in the 1960s. I’ll explain each of these in more detail in the sections to follow. Other useful resources for learning about these four questions include Tinbergen’s original paper (Tinbergen, 1963) and various more-recent review papers (e.g., Bateson & Laland, 2013; Nesse, 2019). For now, be forewarned that none of these questions are easy to answer and all are necessary to develop a complete understanding of a phenomenon.



Figure 1.1: Diagram depicting the hierarchical relationship between Tinbergen’s four types of questions and explanations. Questions and explanations about phylogeny and function are both ultimate, while those about ontogeny and mechanism are proximate.

Ultimate versus Proximate and Tinbergen’s Four

Ultimate questions (or explanations) ask (or explain) why a thing, like a personality trait, exists at all. These questions are typically concerned with the emergence and purpose of a trait in a population of organisms (e.g., a species, taxa, or kingdom) over deep evolutionary time (i.e., hundreds of thousands or millions of years). Ultimate questions are further broken down into Phylogeny and Function within Tinbergen’s framework.

Phylogeny is a fancy word referring to the evolutionary history of a trait. Research on the phylogeny of a trait aims to explain its evolution over deep time from the non-existence of a trait to its current form. Examples of questions about phylogeny include: When did the trait first evolve? Did it evolve independently in different species or do many species have the trait because they share a common ancestor with the trait? Questions about phylogeny are notoriously difficult to answer because they rely on careful comparisons of related species, fossil evidence, and genetics.

Function refers to the evolutionary purpose of a trait. In general, traits become common in a species because they do something to help individuals survive and reproduce (or are intertwined with traits that do). Traits that directly aid in solving the problems of survival and reproduction that organisms face are called adaptations. To make progress on understanding many traits we might be interested in, we can take a gamble and assume that it exists because it somehow solved a problem related to survival and reproduction through evolutionary time.

By trying to identify potential evolutionary functions of a trait, we can develop better questions to potentially learn undiscovered things about the trait and, eventually, to develop explanations for why a given trait works or looks the way it does. Some examples of questions about function include: what problem did this trait help our ancestors solve? How does the design of the trait (e.g., the chemical makeup of stomach acid) help to solve the problem (e.g., digesting food)?

Functional explanations can be difficult to generate because the function of a trait may not be obvious, and in our modern world some traits may not even function the way they are supposed to—this is a phenomenon called *evolutionary mismatch*, which we will explore in a later chapter. Moreover, sometimes it is too easy to come up with satisfying explanations for the evolved functions of traits without really scrutinizing them scientifically; this is called “just-so” storytelling. We must be vigilant to make sure that the functional explanations that we accept have generated novel predictions and supporting evidence through scientific inquiry. (we’ll talk about all this more in a later chapter on evolutionary approaches to human behavior.)

Proximate questions are generally more concerned with identifying how traits work and develop within an individual’s lifespan (or on even shorter timescales). The two types of questions within the proximate level of analysis are questions about *ontogeny* and *mechanism*.

Ontogeny is another fancy word that simply refers to development. Questions and explanations about ontogeny are concerned with the development of a given trait from conception onwards. Ontogeny-focused questions don’t have to focus on the full lifespan, however. Some example questions include: how does the trait change during puberty? How does early life experiences relate to the development of the trait? How does an individual’s genetic code interact with the environment to produce the trait? At what age does the trait emerge or disappear?

Mechanism refers to how a trait works. Mechanistic explanations depend on identifying and describing the physical processes within the body that contribute to the development of the trait. If we can think of organisms as organic robots, then mechanistic questions are designed to assess how the gears, wires, sensors, and software of the robot interact to build the robot and produce its behavior. Some mechanism questions include: How do hormones interact to produce a behavior?

What organs are involved in the producing a behavior? What are the activating rules, or inputs, that trigger the production of the behavior?

Importantly, none of the explanations that we might derive from the four types of questions can provide a complete understanding of a trait by themselves. To fully understand any trait that we might be interested in, it will be necessary to compile explanations from both the proximate and ultimate level of analysis. By organizing research questions and explanations according to this framework we can avoid some common pitfalls that slow our understanding and scientific progress.

Levels of Analysis

To understand personality, we must first understand how the mind-brain works. This is no easy task—humans have been trying to figure this out for hundreds of years! Since the 1950s, thinking of the brain as an information processor has been the guiding principle of cognitive science and has proven quite useful (Miller, 2003). According to this view, the mind-brain can be viewed as an information processing system, just like the computer I’m writing this on. But even simple information processing systems can be pretty difficult to understand. Most of us, including myself, don’t even know how our laptop’s word processor works in much detail... and the mind is much more complex!

So how can we best go about organizing our research questions and explanations to figure out how the mind works in an information processing sense? Lucky for us, the neuroscientist David Marr proposed just such a framework for making sense of complex information processing systems at multiple levels of analysis in the late 1970s.

Marr proposed three levels of analysis: the *computational level*, the *algorithmic level*, and the *implementational level* (c.f., Bechtel & Shagrir, 2015). The three levels are hierarchical, which means that (in general) we must start with the first level and move down to more detailed levels one at a time, especially when it comes to psychological phenomena. If we could develop a good understanding of the mind at each of these three levels, we may be able to more easily figure out where individual differences in thoughts, feelings, and behavior arise.




Figure 1.2: Diagram depicting the hierarchical nature of Marr’s three levels of analysis.

At the **computational level**, we define the problem that a given system is supposed to solve, and ideally why it needs to solve that problem. This is very similar to the functional level of analysis within Tinbergen’s questions: What is the purpose of the thing? Why did it evolve?

Let’s try to define the computational level for a word-processor program. We might say that its purpose is to allow us to get our thoughts out of our heads and onto virtual paper in a way that is easily editable and nicely formatted. We might also include that it should detect spelling errors because we want to share our thoughts with others in a way that can be easily understood.

Now let’s look at a more psychological example. The emotion we call “anger” could be defined at the computational level as a psychological software program, much like a word processor or other computer program. We could say that this anger software’s function is to bargain for better treatment from others because the way others treat us influences our survival and reproduction (Sell et al., 2009). We’ll explore this computational analysis of anger more when we learn about emotions and personality.

At the **Algorithmic level**, we define the steps that the information processing system would need to carry out to solve the problem defined at the computational level. An *algorithm* is a series of steps or rules. The complexity of the algorithm required will depend heavily on the problem that needs to be solved. Moreover, there are likely many potential algorithms that could successfully accomplish the goal and representations (e.g., symbols, concepts) that can be used in the computation.

Let’s go back to our word processor example. Perhaps we could provide an algorithmic description for the automatic spell-checker as follows:

- (1) the user typed a word (e.g., “algorithmic”) into the virtual paper and the computer compares that word to a massive dictionary containing words in the language the user has previously specified
- (2) if an exact match in the dictionary is found, then nothing happens; but if no exact matches are found, the program puts an ugly red line under the word on the screen, looks for words that are very close in spelling to that word, and provides these suggestions to the user.

This is likely an oversimplification of the how spell-check works but it should give you an idea of how we might begin to lay out the steps and rules to define the spell-checker at the algorithmic level. The algorithm is the series of steps and rules, and the representations are the words in the English dictionary (or other dictionary for people writing in other languages).

How about for something more psychological? We can attempt an algorithmic definition for the emotion of anger as follows:

- (1) after an interaction with a person, your brain estimates the value that they place on your wellbeing compared to how much you think they should value your wellbeing;
- (2) if your brain estimates a large discrepancy between their valuation of you and what you think their valuation should be, anger is activated;
- (3) anger produces thoughts, feelings, and behaviors that may be successful in changing how much the person values you in the future (e.g., you make a face so they know you are angry, you threaten to not hang out with them anymore).

Of course, this is an oversimplification of the algorithm(s) underpinning anger. But even this relatively simple algorithmic analysis can allow us to develop more fine-grained descriptions of anger and develop clear predictions (does anger activate when we feel undervalued?). Ultimately, this should help us better understand the emotion and how it works. In this example, the representations used by the algorithm are fairly abstract (VALUE; WELLBEING), and the steps are fairly simple. But we could attempt to make the representations and algorithmic steps more concrete by, for example, using mathematical functions or computer code to model define them. And we can develop more concrete algorithmic descriptions about how the mind might compute value and wellbeing.

At the **implementational level**, we define how the algorithms that solve the problem are set up in the physical world. If the algorithmic level is the software of a program, the implementational level refers to the hardware the program runs on. At this level, we want to point to physical things in the world (e.g., neurons, computer chips, gears, and cogs) that actually carry out the algorithmic computations in our physical reality. This is a very challenging level of analysis to get right because many complex information processing systems, like the mind, are hierarchically organized. So, different processing systems may be implemented in

the same physical space, but they depend on different patterns of interactions of many physical elements.

Let's go back again to our word-processor example. To give an implementational level description we probably have to point to the hardware of the computer. This will likely include many elements. For example, we would need to incorporate the keyboard we use to (mis)type the words in our heads; and the computer chips that convert electrical information sent from the keyboard into symbols on the screen; and the pixels on the screen that display the symbols; and mouse that allows us to right-click the ugly red squiggly line to get the correct spelling for the word we misspelled. And probably much more. I won't pretend to know how all the details are truly implemented on the physical parts of my computer.

And we would hope to be able to do the same for our psychology. If we want to give a description of any psychological program, like anger, at the implementational level, we'd have to point to our brains. But saying that "it's implemented in the brain" is probably too vague to be a useful description of anger because everything about human psychology is ultimately instantiated in the brain!

Ideally, we could point to some more specific part of the brain (e.g., the amygdala or frontal lobe). And eventually maybe to more specific clusters or neurons or interactions among groups of neurons. We might also need to describe how other areas of the body are involved in anger. For example, the muscles and tendons that contort our faces in a way that says, "I'm angry!" without us saying a word. In general, I think we are long way off from understanding the implementational level of many psychological phenomena (like emotions or personality)—especially because we don't yet have a great grasp of things at the algorithmic or even the computational levels of analysis. But maybe you can help change that!

Connecting to Personality

At this point, you may be wondering what any of this has to do with personality. It might not be clear how applying Tinbergen's Four Questions or moving through Marr's three levels of analysis helps us understand a personality trait. To better appreciate how these frameworks can help us study personality, let's practice applying them to personality.

Pick a personality trait to think about for this exercise. I'm going to go with *pessimism* because I'm a bit pessimistic and I like alliteration, so the P in pessimism gives me some alliteration with my first name and it also describes me: *Pessimistic Patrick*. You should think about a different trait here—maybe one that describes your own personality and makes some nice alliteration with your first name.

How might we apply Tinbergen's Four Questions to [pessimism / your trait]? Let's start with the ultimate questions concerning function and phylogeny.

For **function**, we want to ask questions and seek answers that allow us to understand why the trait evolved. We could also ask questions like: What is the purpose of [pessimism]? How might [pessimism] have helped our ancestors survive and reproduce? How might traits that aided our ancestors' survival and reproduction create patterns of thought, feeling, and behavior that we label as [pessimism]¹. Insert your own trait term into the brackets to see how to ask functional questions for your trait.

I could tentatively argue that the evolutionary purpose of pessimism is to prevent individuals from engaging in activities that are unlikely to succeed, which may prevent unnecessary energy expenditure. Or perhaps the purpose of pessimism is to magnify the worst in things so that they can be more easily identified and improved, which may result in things improving over time. Either of these will conjectures² will do for now. What potential purposes can you think of for the trait you are analyzing?

For **phylogeny**, we want to ask questions and find answers that help us understand how the trait evolved over time. For example, we could ask: When did pessimism first arise in human evolution history? Do any non-human animals exhibit a trait like pessimism?

These kinds of questions can't be answered very easily, but based on our function assessment above, we could make some educated guesses about phylogeny. If pessimism's purpose is indeed to prevent engagement in activities that are unlikely to succeed, we might expect that the trait of pessimism—or something like it—could go all the way back to a common ancestor that first had to assess the likelihood of success of potential activities. We would need vast amounts of data from many sources to actually test this. For now, it's just sort of fun to think about some ancient common ancestor of ours experiencing pessimism. What kinds of phylogeny questions can you ask about your trait?

Now let's turn to the proximate questions, ontogeny and mechanism.

For **ontogeny**, we want to ask questions that let us develop an understanding of the development of the trait. Questions like “What age does pessimism appear?” and “Do people get more pessimistic as they age?” might be good places to start. We could then design a research study to provide some insight into these questions. What ontogeny questions can you make about your trait? Try to come up with some different ones than I wrote above.

For **mechanism**, we want to ask questions that will help us explain how the trait works within an individual. Some potential starting questions might be, “what sorts of situational factors are associated with pessimism?” Again, we should then be able to start designing research studies to clarify these questions. Can you develop some questions about mechanism for your trait?

¹This question is essentially assuming that pessimism is a non-functional *byproduct* of other traits that do have a function. We'll learn more about this in future chapters.

²Since I don't have any real data to back either of these claims, they are currently “just-so” stories (or, if you're feeling generous, they're hypotheses). But if I generate and test some novel predictions, I could potentially see which story about the function of pessimism seems more likely, or which gives the most new insight about pessimism.

By developing some questions for each of Tinbergen's four areas, we have begun to outline clear directions for a research program that will deliver a holistic understanding of the trait we are interested in. No single researcher must directly study all of these four areas. Individuals or groups of researchers can focus on one or two they find most interesting or practical to study. But we should keep in mind that all four areas are important and we don't want to end up neglecting any set of questions. Which of Tinbergen's Four Questions are you most interested in?

Now let's continue our analysis by going through Marr's three levels. I'll continue with pessimism, and you can continue with your chosen trait term.

We already implicitly developed two potential **computational** descriptions of pessimism when we made hypotheses about the function of pessimism. I'll restate one potential function here: pessimism prevents individuals from engaging in activities that are unlikely to succeed in order to save energy. This implies that the function of pessimism is energy conservation.

Now that we have a tentative computational description, we can try to work out an **algorithmic** description. Let's try to use an equation to represent the algorithmic computation that defines how much a person feels pessimism about a potential activity. The algorithm could work as follows:

1. estimate the energy to be gained by a potential activity. We'll call this estimate (E_{gain}) and we'll use whole numbers 0-100 to represent the energy;
2. estimate the minimum energy that would need to be spent on the activity. We'll call this estimate (E_{cost}). We'll use whole numbers from 0-100 to represent energy;
3. estimate the odds of success of the activity. Let's call this estimate O . This is a probability, so we'll use a decimal value ranging from 0-1 to represent this probability;
4. Compute a ratio of the minimum estimated energy cost compared to the potential energy gain scaled by the odds of success: $P = \frac{E_{cost}}{O \times E_{gain}}$
5. Activate pessimism (P) in direct proportion to the value estimated in step (4)
6. Produce thoughts (e.g., "This will never work"), feelings (e.g., negative affect, lethargy), and behavior (e.g., withdraw effort, voicing doubts) that effectively reduce or prevent energy expenditure on the potential activity.

This algorithm makes it so that the value representing the amount of pessimism a person feels (i.e., P) will be higher when as (a) the energy cost increases, (b) the odds of success decrease, and (c) the potential energy gain decreases. Conversely, pessimism will be lower when the energy cost is low, when the odds of success are high, and when the potential energy gain is high. You can see this for yourself by plugging in different values into the equation in step 4 using whole numbers 0-100 for E_{cost} and E_{gain} and a decimal ranging from 0-1 for O to calculate P .

Now that we have this algorithmic description, we could develop empirical tests to see how well human pessimism feelings can be predicted by this algorithm. For example, ask people how much pessimism they feel towards a range of situations that systematically vary in energy costs, energy gains, and odds of success to see how those parameters interact to predict self-reports of pessimism.

Keep in mind that this algorithm (and probably any algorithmic description of human behavior) is just a useful fiction to help us describe clearly how we think pessimism—at least as we’ve defined it in the computational description—works. If people are imagining a different computational function of pessimism, perhaps more like the alternative function I came up with in the Tinbergen analysis above, then this algorithm will likely not be a good or useful description.

The last of Marr’s levels is the **implementational level**. Of course, any psychological phenomenon will ultimately need to be linked to brain activity to develop an implementational level explanation. We’d ideally want to be able to point to regions of the brain, networks of neurons, and specific neurotransmitters to understand how the algorithmic computations underpinning pessimism work. And we can’t really do that without a good understanding of the algorithmic and functional levels. I’m no neuroscientist, so I won’t speculate about the implementational level here other than that it’s got to be implemented by our neurons and neurotransmitters.

Ultimately, by developing description of pessimism—or any trait—at the computational, algorithmic, and implementational levels of analysis, we can begin to provide more-complete and detailed explanations of psychological phenomena. In theory, this should allow us to recreate the phenomenon of interest, using different implementational components than our biology uses. For example, we could potentially simulate the trait using the software and hardware of a computer or robot instead of the software and hardware of organic organisms.

Wrapping Up and Building on These Foundations

In this chapter, we’ve just begun our journey to understand the complexities of personality. It’s understandable that the vast field of psychology can initially appear bewildering, with its disconnected findings, loosely defined theories, and myriad empirical questions. But as we’ve seen, there are foundational frameworks at our disposal to make sense of this seemingly chaotic landscape. We’ll rely on these frameworks frequently throughout this course.

By applying Tinbergen’s Four Questions, we gain insight into the ultimate and proximate aspects of personality traits, dissecting their evolutionary history and functional purposes. We’ve explored how these questions can help us understand the “why” and “how” behind personality traits, uncovering their potential adaptive functions and development across an individual’s lifespan.

Marr's three levels of analysis, from the computational to the algorithmic and implementational, provide us with a structured approach to dissecting the inner workings of personality traits. We've seen how these levels allow us to create hypotheses and models for understanding traits like pessimism in a systematic manner.

The pursuit of a complete understanding of personality is no easy task, and it requires collaboration among researchers specializing in different levels and areas of analysis. By leveraging these frameworks, we can begin assembling the pieces of the puzzle, one by one, to put together a more comprehensive picture of personality and human psychology more broadly. In the chapters to come, we'll continue relying on these foundational frameworks and explore how they can be applied to shed light on various aspects of human psychology and personality.

References

- Anvari, F., & Lakens, D. (2018). The replicability crisis and public trust in psychological science. *Comprehensive Results in Social Psychology*, 3(3), 266-286.
- Bateson, P., & Laland, K. N. (2013). Tinbergen's four questions: an appreciation and an update. *Trends in ecology & evolution*, 28(12), 712-718.
- Bechtel, W., & Shagrir, O. (2015). The non-redundant contributions of Marr's three levels of analysis for explaining information-processing mechanisms. *Topics in Cognitive Science*, 7(2), 312-322.
- Eronen, M. I., & Bringmann, L. F. (2021). The theory crisis in psychology: How to move forward. *Perspectives on Psychological Science*, 16(4), 779-788.
- Miller, G. A. (2003). The cognitive revolution: a historical perspective. *Trends in cognitive sciences*, 7(3), 141-144.
- Nesse, R. M. (2019). Tinbergen's four questions: Two proximate, two evolutionary. *Evolution, Medicine, and Public Health*, 2019(1), 2-2.
- Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift für tierpsychologie*, 20(4), 410-433.

2 Past Perspectives

A thorough and systematic understanding of personality will likely require the scientific method. But humans have only been applying something like the scientific method to understand personality and individual differences for a hundred years or so. So, for most of history, ideas about personality have been largely nonscientific. What can we learn from these past perspectives on personality that can give us a better appreciation for scientific approaches to understanding personality?

For one, we might just find it interesting to know how long humans have been trying to understand personality and individual differences. If humans around the world thought about personality for a long time, perhaps this suggests that personality differences have been important to humans for a long time. On the other hand, if caring about personality is a more recent phenomenon, maybe personality differences are only an important factor in the contemporary world for some reason.

Second, getting a sense of how long people have been thinking about personality can help us avoid hindsight bias, which is the feeling that things are obvious or easy after we've already learned about them or finished the task. By learning about past perspectives on personality, we can appreciate just how long people have been trying to understand personality—arguably without making much progress. When we see what past perspectives on personality generally were, it will be clear that things we learn about personality are not in fact obvious.

Third, examining how long people have been trying to understand personality and seeing how little progress has been made can help us appreciate that studying personality, or really any aspect of human psychology, is hard! We humans have unlocked many secrets of the universe: we've sent people into space, we've split the atom, we've created vaccines to nullify many dangerous diseases, we've created machines to talk to people anywhere in the world; but we can barely begin to understand the personalities and psychology of the humans that achieved these feats. Personality is largely still a mystery.

So, how long have humans been thinking about personality? What were some perspectives on personality in the past? Let's explore.

Is Personality Ancient?

Personality in the Oldest Written Story

The oldest written story we have evidence for is the *Epic of Gilgamesh*. It is a Mesopotamian story found written on stone tablets that date back to around 2200 BCE — that's over 4000 years ago! And the story itself was probably transmitted verbally for a long time before it was written down.

The story is about an ancient king, Gilgamesh. In the beginning of the story, Gilgamesh is generally making life miserable for his subjects. The gods create a rival, Enkidu, to stop Gilgamesh. But Gilgamesh ends up beating Enkidu when they fight. So, naturally, they become best buds. Gilgamesh and Enkidu have many epic adventures until Enkidu dies. Gilgamesh then spends the rest of the story grieving and obsessing over his own mortality. He goes through a series of trials and tribulations in search of a way to defeat death, before ultimately making peace with his mortality.

The obvious takeaway from the story is that death is inevitable, and we should enjoy life while we can. But we could also view it as a story about personality change. Gilgamesh's behavior and feelings towards others changes quite a bit from the beginning to end: he started out a terrible person trying to control everyone and everything, and in the end he makes peace with his own mortality. Maybe the real epic of Gilgamesh is the personality change along the way.

But were the ancient humans who wrote down the epic including information about the characters personality? Du et al. (2023) examined the Epic of Gilgamesh to see what personality relevant information it contained and how that information might have changed as it was transcribed over the course of 1000 or so years. They found that even the earliest version of the story contained person descriptor words that are relevant to personality. Moreover, the number of personality descriptor words increased over time in different versions of the story. Thus, it appears that even 2000 years ago, humans wanted to know about others' personality.

What About Earlier?

The Epic of Gilgamesh is old but not super old. Humans have existed for at least 200,000 years. Did super ancient humans care about individual differences and have ideas about personality? This is a difficult question to answer because we can't look back in time to see what they thought. But I think we can be fairly certain that our ancestors cared about personality for several reasons. Below I outline some reasons to think that personality differences existed and were important enough to care about throughout deep in our evolutionary past.

As far as we know, ancient humans were not clones, so they probably would not have been identical in every way. Individual differences in cognition and behavior would almost certainly

have existed. Ancient humans had to make lots of decisions about others that would have impacted their survival and reproductive success: who is a good mate? Who is a good hunting partner? Who is a reliable cooperator? Who is a good worker? Knowing about differences in cognition and behavior between people is information that would be useful for our ancestors to determine the answers to those questions. Because of the way natural selection works (which we'll unpack in a later chapter), it would be surprising if our brains didn't evolve to detect differences between people that are relevant to solving problems of survival and reproduction.

As we'll see in the next chapter, personality as we know it depends largely on language. Language is probably at least 100,000 years old. Once language evolved, it would have become possible, at least in principle, to share information about personality attributes of others. Moreover, as language became more refined, our ability to communicate and think about the small differences between people would have become more refined as well. Thus, language and personality concepts could have coevolved culturally in a feedback loop: more fine-grained language could make it possible to identify small differences, and identifying small differences could lead to more fine-grained language concepts.

Of course, this is all a "just-so-story". But I can't really come up with a good reason why our ancestors wouldn't have paid attention to personality, and why personality concepts wouldn't have emerged in language fairly early on. Ultimately, I think it's pretty safe to say that humans have been thinking about personality and individual differences, at least implicitly or unconsciously, for a very very very long time. But what can we know about what these ancient humans thought about personality?

A Nonexhaustive History of Historical Personality Perspectives

A Note on Missing Data

Personality doesn't fossilize. We can only learn about what ancient humans thought about personality if those ancient humans happened to record their thoughts on material that would last long enough for us to find it. The overwhelming majority of personality-relevant information about cultures from around the world are lost to time.

So, we don't have data on what most cultures thought of personality. For example, we don't have good evidence for what indigenous Californians thought about personality in the thousands of years before colonization. But absence of evidence is not evidence of absence.

It is important to appreciate that the cultures we do have data from may not be representative of all human cultures that have ever existed. Still, by looking at the data we do have, perhaps we can start to get a sense of what sorts of things ancient humans might have thought about personality.

Ideal Personalities

Ancient Confucian writings going back to 500 BCE discuss the concept of *Junzi*, which could be thought of as the closest any individual person could get to the ideal personality. Junzi has three components: *Ren* (benevolence), *Zhi* (wisdom), and *Yong* (courage). Under this view, people could cultivate different levels of each component and each component is causally related to one's behavior and thought. In contrast to Junzi, Confucian writings also discuss a personality concept of small-mindedness or pettiness, called *xiaoren*.

The Bhagavad-Gita, an ancient Hindu text going back to around 200 BCE talks about personality being composed of three “modes”, the *Tri Gunas*. The three modes are: *sattva* (holiness), *rajas* (desire), and *tamas* (darkness). In this belief system, a person's actions are determined by the dominant mode. The relative dominance of each mode could be attributed to many things: mystic formulas, occupation, birth, ancestry, thought, etc.

Humorous Humorism

Humorism is the belief that illnesses and differences in behavior correspond to the balance of four fluids (“humors”) in the body: yellow bile in the liver; black bile from the kidneys; phlegm from the lungs; and blood from the heart. This belief system is often attributed to the ancient Greek physician-philosophers Hippocrates and Galen. But beliefs like humorism may actually go back further to ancient Mesopotamia or Egypt.

Under this view, each of the four humors corresponded to a different “temperament”, or characteristic patterns of behavior—much like a personality trait. Too much yellow bile supposedly leads to a choleric temperament, which is passionate, brave, and determined. Too much black bile leads to a melancholic temperament, which consists of being detached, worried, or unhappy. Too much blood led to sanguine temperament, which tends to be joyful, eager, and optimistic. And too much phlegm led to a phlegmatic temperament of peace, dependability, and thoughtfulness. So, next time you feel sad, just find a way to increase your blood-to-black-bile ratio.

These beliefs might sound humorous to us now, but humorism was prevalent for over 1,000 years and continued to be popular through the Middle Ages. It wasn't until the theory of germs was developed that humorism's influence on medicine was remedied. This is good opportunity to reflect with humility on our own popular theories and beliefs; there could be many popular ideas about personality today that could seem as ridiculous to future generations as humorism does to us.

Freud and the Neo-Freudians

Freud's Model of the Mind

Freud is probably the one personality theorist most people could name off the top of their head, which is a little strange given how little of his specific theory has survived contact with evidence. His influence and his evidence are two different things, and it helps to keep them apart.

Freud's central idea was that much of mental life happens outside of awareness, and that personality comes out of conflict between motives we cannot fully see. He described that conflict using three parts of the mind. The **id** holds our most primitive urges—hunger, thirst, sex, aggression—and wants them satisfied immediately, whatever the consequences. The **superego** is the set of social rules we absorb from parents and culture; it works as a conscience, producing pride when we meet its standards and guilt when we fall short. The **ego** sits between the two, trying to get the id what it wants in ways the world will actually tolerate.

A healthy personality, on this account, is one where the ego manages that negotiation well. When it doesn't, Freud thought the result was anxiety and neurosis. A person ruled by the id would be impulsive and self-centered; a person ruled by the superego would be rigid and guilt-ridden, denying themselves even ordinary pleasures. Freud also proposed that the ego protects itself with **defense mechanisms**: unconscious tactics like repression, which pushes threatening thoughts out of awareness, or projection, which pins our own unacceptable impulses on somebody else.

Underneath all of it, Freud placed the sex drive, or *libido*, as the main engine of human motivation. That assumption did an enormous amount of work in his theory, and it is the first thing his followers threw out.

Psychosexual Stages

Freud also argued that personality is largely settled in early childhood, as the id's pleasure-seeking moves through a fixed sequence of bodily focuses. If a child's needs at a given stage were handled badly—too much indulgence or too little—the person would stay *fixated* there, and it would show up in their adult personality.

The stages are the **oral** stage (birth to about 1 year), centered on the mouth and feeding; the **anal** stage (about 1 to 3), centered on bowel control and toilet training; the **phallic** stage (about 3 to 6), centered on the genitals and on the child's relationship with each parent; a **latency** period from about 6 to puberty, when Freud thought sexual interest goes quiet; and the **genital** stage from puberty onward, when those interests return and get directed outside the family.

The phallic stage carried the most theoretical weight. Freud proposed that boys desire their mother and treat their father as a rival—the **Oedipus complex**—and resolve the problem by identifying with the father rather than competing with him. His account of girls ran through *penis envy*: the idea that a girl registers the anatomical difference as something she lacks, and that this shapes her development. The term *Electra complex* is often used for the female parallel and just as often credited to Freud, but it was Jung’s coinage, and Freud never took it up.

All of this rested on a very narrow evidence base. Freud worked from his own recollections and from his interpretations of what his patients told him about their lives and dreams, and he treated that as enough to establish that these stages were universal features of childhood. By any modern standard it isn’t. Introspection and clinical impression are a perfectly reasonable way to come up with a hypothesis; they are not a way to test one, and Freud mostly did not test his.

So why does Freud still get a section? Because the general shape of his thinking outlasted the specifics. The ideas that mental life is partly unconscious, that people are pulled by motives that conflict with one another, and that early experience matters for later personality are all still live in personality psychology, usually in forms Freud would not recognize, and resting on evidence he never collected.

Neo-Freudians

Freud’s ideas were influential largely through the disagreements they provoked. The psychologists who followed him, usually called neo-Freudians, generally accepted that unconscious motives matter and that development shapes personality, but rejected his insistence that sex was the engine driving all of it. Three are worth knowing.

Alfred Adler: Inferiority and Striving

Adler swapped Freud’s master motive for a different one: the drive to be superior to others, or at least to avoid being inferior to them. On this view, much of what we do is compensation for the places we feel we fall short.

Adler also proposed that birth order shapes personality. Oldest children, having lost their parents’ undivided attention, become overachievers; youngest children are indulged; middle children are left to work out something in between. The idea is memorable and has been repeated for a century. It has not held up especially well. Research has provided little support for systematic birth-order effects on personality.

Karen Horney: Anxiety and Security

Horney, one of the first women trained as a psychoanalyst, went straight at Freud's claim that girls envy boys' anatomy. Any such envy, she argued, is far more plausibly about privilege than about anatomy: in a patriarchal society, boys simply get more. She added that one could just as easily claim men envy women's ability to give birth, which is a fairly efficient demonstration of how much these claims depend on what the theorist assumed going in.

Her own account put anxiety, rather than sex, at the center of personality development. Children whose needs go unmet, who experience loneliness or isolation, develop a basic anxiety they then have to manage, and they seek belonging and security in response. Horney described three ways of doing it: moving toward people, moving away from people, and moving against them. All three are ordinary strategies that everyone uses. The problem, in her account, is rigidity: leaning on one of them regardless of what the situation actually calls for.

Carl Jung: The Collective Unconscious and Introversion

Jung broke with Freud on the same point as the others—he did not accept that sex was the primary motivator—and also on the scope of the unconscious. Freud's unconscious was personal, assembled out of an individual's own history. Jung proposed that underneath it sits a **collective unconscious**: inherited patterns, which he called **archetypes**, shared by everyone. He pointed to recurring figures in the myths and folklore of unconnected cultures—the hero, the maiden, the sage, the trickster—as evidence that the same themes are handed down biologically. It is a striking idea that has attracted very little empirical support, and the inference is shaky: cultures separated by an ocean can end up telling similar stories for reasons much simpler than inherited memory.

Jung's more durable contribution was far less exotic. He described two orientations toward the world, **extraversion** and **introversion**, distinguished by whether a person's energy comes from engaging with other people or from their own inner activity. Some version of this distinction shows up in nearly every modern model of personality, which arguably makes it the most successful idea to come out of this whole tradition. He also described the *persona*, the face we present in public, which he treated as a compromise between who we actually are and what other people expect us to be.

Jung's types are also the reason for one of the most widely used personality tests in the world. The *Myers-Briggs Type Indicator* (MBTI) was built loosely on his ideas, sorting people on introversion versus extraversion, thinking versus feeling, intuition versus sensation, and judging versus perceiving. It is popular on dating websites and in workplaces, and its central problem is structural rather than a matter of the individual questions. The MBTI reports a *typology*, sorting people into discrete categories, but the traits underneath it are continuously distributed. Any set of categories therefore has to cut arbitrary lines through the middle of real distributions, which means two people sitting just either side of a cutoff get labelled as different types despite being nearly identical—and someone close to a boundary can come out

as a different type on a retest without their personality having changed. Measuring where a person falls along a continuum, rather than which box they belong in, has a much stronger psychometric case. These are concepts we will learn more about in a later chapter.

Conclusion

Humans have probably cared about personality for very long time. And we've been trying to understand personality for at least a few thousand years. However, as this chapter shows, many historical approaches to understanding personality have been non-scientific. Most of the historical beliefs and so-called theories about personality were based on self-reflection and opinions about what makes a healthy or virtuous person. But this approach to understanding personality leaves out a lot of the variation in human behavior that makes up the many diverse personalities that exist. The scientific approach to understanding personality that we will explore throughout the rest of the chapters brings large amounts of data to allow us to systematically study personality, while hopefully making it harder for our own subjective biases and beliefs to creep in.

References

Du, A. H., Karl, J. A., Fetvadjiev, V., Luczak-Roesch, M., Pirngruber, R., & Fischer, R. (2023). Tracing the evolution of personality cognition in early human civilisations: A computational analysis of the Gilgamesh epic. *European Journal of Personality*.

Earlier versions of the sections on Freud and the neo-Freudians were adapted from the [Psychology, Personality materials on OER Commons](#), licensed under [CC BY 4.0](#). No authors are listed for those materials. Those sections have since been rewritten.

3 Contemporary Perspectives

Too Many Theories, Not Enough Data

In the last chapter, we learned about the many varied perspectives on personality throughout history. These so-called “theories” were little more than hunches or guesses that people made based on their own experiences with people (e.g., Freud’s therapy patients). But individual experiences are severely limited, and often biased. We should not expect that a theory of personality based on any single individual’s personal experiences and perceptions will help us understand personality. What we need is a way to pull together many experiences and perceptions in a systematic way. That’s where the analytic, data-driven approach to personality comes in.

By collecting data about personality from large groups of people in the populations we aim to understand, we are much more likely to be able to begin to systematically identify and categorize personality traits and processes. In this chapter, we’ll examine how the data-driven approach to understanding personality has helped the field of personality science advance over the last 100 or so years. We’ll explore the process that led to the current dominant model(s) of personality.

The Dictionary as Data

A major goal of personality researchers throughout much of the 1900s was to create a taxonomy of personality traits. A taxonomy allows for systematic identification and categorization of things—whether those things are plants and animals, cars, learning styles, or personality traits. A personality taxonomy would allow researchers to organize and understand how the many thousands of ways that humans are similar and different from one-another in tendencies of thoughts, feelings, and behavior. But with so many potential personality traits where do we even start?

Modern personality research starts in the dictionary. In everyday language, people use words to describe the important differences between people in their social world. Personality theorists reasoned that words that capture important differences about people and their personality traits should therefore become encoded in our lexicon over time. Dictionaries contain words commonly used in a given language, which makes dictionaries a kind of passively data-driven record of many peoples’ experiences and perceptions. So, perhaps we can look in the dictionary

to get a sense of how many personality traits there are! This line of reasoning—that personality traits are encoded in our language and that trait terms can be used to learn about personality—is called the lexical approach to personality.

Allport's Lexical Analysis

Gordon Allport was one of the first researchers to systematically use the dictionary to identify potential personality traits (Feeney, 2020). Allport and his colleague, Henry Odbert, went into the dictionary and recorded ~18,000 words that described human personality and behavior. From this list, they organized words into four categories. The first category identified ~4500 words that describe dispositions (e.g., outgoing). The remaining categories contained words describing behavioral states (e.g., laughing), social evaluations (e.g., high-status), physical traits (e.g., tall), or abilities (e.g., athletic), which Allport did not consider to be personality traits per se.

It seems that over 4000 words would be a pretty good starting point from which to try to systematically organize personality traits. But many of these words are probably synonymous with each other (i.e., they have the same meaning), making them somewhat redundant. We want to figure out how many unique personality traits are out there in real human populations, not just language. This is where collecting data from people becomes necessary.

People as Data

To get a better sense of how many personality traits are being described by 4000 or so trait terms in the dictionary, we need to get a sense of how their use overlaps in the population. That is, we need to be able to see whether people who are commonly described by a given trait adjective term (e.g., outgoing) also tend to be described using other adjectives (e.g., fun-loving, personable). To put it another way, we want to be able to group trait terms together within a taxonomy based on how they overlap—how they are correlated—in the population. The next stage of researchers in the lexical tradition did exactly this to develop taxonomies about personality, which ultimately led to the dominant models of personality structure in contemporary personality research.

Cattell's Factor Analyses and 16 Factors

Raymond Cattell was the first psychologist to attempt to use a data-driven process to determine how to group the many words that describe personality into more manageable set of clusters (i.e., groups of words that describe similar personality traits). First, Cattell reduced the set of words that Allport and Odbert made from ~4000 down to 171 by eliminating words that he thought were simply synonyms. He then conducted several studies where he gave the

list of words to people and had them rate the extent to which each word was true of them. The ratings provided by these subjects allowed Cattell to compute the correlations between each pair of trait terms, resulting in 14,335 correlations. Computers were not yet available for these computations, so Cattell and his assistants had to lay out these correlations on paper on the ground to organize them—they took up 14 square feet!

Cattell used a statistical method called factor analysis, which he developed along with his advisor, Karl Pearson (creator of the correlation coefficient) to organize the trait terms. A full technical explanation of factor analysis is beyond the scope of this chapter, but it will be important to develop a conceptual understanding of how it works. Simply put, it takes the correlations among variables—in this case traits—and groups based on how strongly correlated they are. For example, if people who rate themselves as outgoing also tend to rate themselves as assertive, sociable, and fun-loving (i.e., these traits are highly correlated), then factor analysis groups them together. These groups are referred to as “factors” in factor analysis literature. The factors are then labelled by the researcher according to their subjective assessment of the content of each cluster (e.g., Extraversion). If most traits are highly correlated, then there will be fewer factors suggested from the factor analysis results; conversely, if most traits are weakly correlated there will be more factors suggested.

Ultimately, after several studies and some somewhat arbitrary revisions and additions, Cattell ended up finding evidence for 16 personality factors. He then developed questionnaires to measure his 16-factor model of personality (i.e., 16PF). But, as we will soon see, the 16-factor model was probably a mistake.

From 16 Factors to the Big Five

Following Cattell’s analyses, many other researchers used factor analytic methods to examine how many personality factors there might be (e.g., Fiske, 1949; Tupes & Christal, 1961; Norman, 1963; Digman & Takemoto-Chock, 1981). But they could not reproduce the 16 factors. Science depends on being able to reliably produce a set of results following a set of procedures. So, if the 16 factors could not be reliably reproduced, we can’t include them in our scientific understanding of personality structure. Some subsequent investigations of Cattell’s methods and analyses suggest that Cattell made a calculation error at some point that led him to the wrong conclusion—which wouldn’t be surprising since all this had to be done by hand before computers. However, we can’t know for sure what went wrong.

Even though later researchers couldn’t find Cattell’s 16 factors in their own analyses, a pattern did consistently emerge across studies. These different teams of researchers each tended to find that five factors consistently emerged from the trait rating data. These five factors, which came to be known as the **Big Five**, were typically labelled as surgency, agreeableness, dependability, emotional stability, and culture/intellect. In the following decades, the Big Five factors were morphed and refined through additional studies and debates. The most popular

current formulations and labels of the Big Five traits are Extraversion, Agreeableness, Conscientiousness, Neuroticism (or Emotional Stability), and Openness (Diener & Lucas, 2023). People often use a mnemonic like CANOE or OCEAN to remember the Big Five traits based on their first letters.

Each of the Big Five factors represents a continuous distribution, meaning that people differ in the quantitative levels of each trait within a population. Under this model, everyone has some level of each trait—this taxonomy based on categorical or discrete typologies. For example, a given person might be average on Extraversion, be above average on Agreeableness, average on Conscientiousness, be well below average on Neuroticism, and very low on Openness, when compared to others in the population; another person could have very different levels. Under this framework, researchers don't generally refer to people to as "Extraverts" or "Neurotics", but rather refer to levels of Extraversion and Neuroticism (e.g., highly Extraverted people, people low on Neuroticism).

Each factor has a specific description which we will explore below. In some models and measurement scales for the Big Five, the factors are also hierarchically organized so that each factor has smaller subgroups within it. Under the Five Factor Model, for example, each of the five factors has six subgroups called facets. Below, I describe each Big Five factor and its facets.

Extraversion

Extraversion refers to a person's sociability, assertiveness, and overall level of energy in social interactions. Those high in extraversion tend to be outgoing, talkative, and enjoy socializing. Those lower on extraversion—often conversationally referred to by laypeople as introverts—may prefer quieter, more solitary activities.

The facets of extraversion are Warmth, Gregariousness, Assertiveness, Activity, Excitement Seeking, and Positive Emotionality. Warmth is characterized by friendliness and the ability to make friends easily. Gregarious individuals enjoy being around people and are often the life of social gatherings. Assertiveness reflects a person's tendency to take charge of social situations. Activity is characterized by a preference for staying busy and always having something to do. Excitement Seeking is marked by a desire for social stimulation and sometimes a willingness to take risks in the pursuit of excitement. Positive emotionality is associated with general cheerfulness and optimism.

Conscientiousness

Conscientiousness describes an individual's degree of organization, detail-orientation, responsibility, and resoluteness in striving toward objectives. Those ranking high in conscientiousness typically demonstrate reliability, meticulous attention to detail, and unwavering dedication

in the pursuit of their aims. Conversely, those with lower conscientiousness tend to be more spontaneous and exhibit a lower tendency for meticulous planning and orderliness.

The facets of Conscientiousness are Competence, Order, Dutifulness, Achievement Striving, Self-discipline, and Deliberation. Competence reflects a desire to excel and achieve high standards in tasks and responsibilities. Order is characterized by a preference for neatness, organization, and tidiness. Dutifulness involves a strong sense of responsibility and reliability. Achievement Striving represents a drive to set and accomplish challenging goals. Self-Discipline is the ability to control impulses and stay focused on tasks. Deliberation is the tendency to carefully consider decisions before acting.

Agreeableness

Agreeableness gauges a person's interpersonal demeanor and relationships with others. Those scoring high in agreeableness are typified by their friendliness, empathic disposition, and predisposition toward cooperation. Conversely, individuals with lower agreeableness may display a heightened competitive spirit or a proclivity for skepticism regarding the intentions of others.

The facets of Agreeableness are Trust, Straightforwardness, Altruism, Compliance, Modesty, and Tendermindedness. Trust involves a belief in the goodness and reliability of others. Straightforwardness is characterized by honesty and a tendency to communicate directly. Altruism reflects a selfless concern for the well-being of others. Compliance involves a willingness to conform and cooperate with others' wishes. Modesty is the tendency to downplay one's own achievements and importance. Tendermindedness represents a sensitivity to the feelings of others and a caring attitude.

Neuroticism (Emotional Stability)

Neuroticism, or conversely, emotional stability, represents an individual's emotional resilience and stability amidst stressors and negative situations. High neuroticism individuals are susceptible to heightened levels of anxiety, mood volatility, and may manifest heightened reactivity to stress-inducing stimuli. In contrast, individuals low on neuroticism tend to exude a more tranquil affective disposition and demonstrate superior emotional management capabilities.

The facets of Neuroticism are Anxiety, Hostility, Depression, Self-Consciousness, Impulsiveness, and Vulnerability. Anxiety involves a tendency to worry, feel anxious, and experience nervousness. Hostility reflects a disposition towards anger, irritability, and frustration. Depression represents a susceptibility to feelings of sadness and hopelessness. Self-Consciousness is characterized by a heightened awareness of self and fears of being judged. Impulsiveness involves acting on urges and desires without thinking through consequences. Vulnerability is the degree to which one is sensitive to stress and prone to emotional fluctuations.

Openness

Openness signifies an individual's propensity to engage with innovative and unconventional concepts, encounters, and emotions. High scorers in this trait are characterized by their proclivity for imagination, intellectual curiosity, and an eagerness to embrace novel experiences. In contrast, individuals scoring lower in openness tend to exhibit a preference for established routines and the comfort of familiarity in their cognitive and experiential domains.

The facets of Openness are Fantasy, Aesthetics, Feelings, Actions, Ideas, and Liberalism. Fantasy is characterized by a rich inner life, creativity, and a penchant for artistic expression. Aesthetics involves an appreciation for art, music, and aesthetics. Feelings is characterized by a deep emotional engagement with the world, often experiencing intense feelings. Actions reflects a willingness to explore the unknown and try new things. Ideas involves a desire for knowledge, complex or abstract problems, and intellectual pursuits. Liberalism represents a preference for progressive and unconventional beliefs.

The Significance of the Big Five

The emergence and widespread adoption of the Big Five factors as a fundamental framework for comprehending human personality has played a pivotal role in the field of personality psychology. Throughout most of the 1900s, scholars engaged in protracted debates about how to categorize the myriad ways in which individuals differ and resemble one another. The consistent recurrence of the Big Five structure in numerous attempts to organize personality variation offered compelling evidence that it was capturing something essential about the structure of personality.

As the Big Five gained increasing acceptance as a foundational framework for organizing personality traits, particularly during the 1990s and early 2000s, it marked a turning point in the field. This acceptance allowed researchers to move beyond the initial struggle of deciding which traits to investigate and how to arrange them. Instead, they could direct their collective focus on the Big Five personality traits themselves and build up our understanding of these personality dimensions over time. In essence, the Big Five provided a general structure to describe personality, freeing researchers to make substantial progress in other goals of personality science.

A Sixth Personality Factor?

In the early 2000s, researchers Kibeom Lee and Michael C. Ashton conducted their own lexical analysis studies of personality that included more terms related to ethicality and morality (e.g., greed, fairness, honesty), which were largely excluded from the earlier lexical analyses stemming from Allport's word lists. By including these terms and conducting studies

in a wide range of cultures, Lee and Ashton found that six factors were the most appropriate way to describe the structure of human personality variation. For the most part, the factors that they found largely matched with the Big Five factors (albeit with some small differences): **E**motionality (corresponding to Neuroticism), **eX**traversion, **A**greeableness, **C**onscientiousness, and **O**penness to Experience. However, they found one important dimension of personality variation was seemingly missing from the Big Five framework: the factor of **H**onesty-Humility. Together these six factors make up the HEXACO model of personality, which is also growing in popularity as a successor of the Big Five model.

As you might guess based on the extra trait terms included in the analyses, the Honesty-Humility dimension focuses on aspects of personality related to ethical behavior, sincerity, and interpersonal trust. Understanding this dimension can provide insight into how individuals approach moral and ethical choices, as well as how honest and humble they are. I would think this is important information about someone's personality that we would want to know for many types of relationships (e.g., business partner, romantic interests, friends), so it is surprising that it remained essentially "hidden" to personality research for so long.

So How Many Personality Dimensions Are There?

In our quest to understand human personality, it's important to recognize that no single model can claim absolute truth or objectivity in summarizing the entirety of personality variation. These models, whether they propose five or six dimensions, are fundamentally shaped by subjectivity at various stages of their development.

Firstly, subjectivity is inherent in the process of data collection used to inform factor analyses. The terms and measurements employed to capture aspects of personality inevitably rely on individual interpretation and perspective. The subjective judgments about which trait terms to include, for example, influence what data can even inform the statistical tests we run. If we don't put the right terms in, we won't get the right results.

Moreover, subjectivity extends to the interpretation of results derived from factor analyses. Researchers must exercise judgment and discretion when deriving meaning from factor analyses. For example, researchers have to decide how many factors should be retained and there are potentially infinitely many ways of interpreting factor analytic results depending on specifics of the analyses (e.g., the factor rotation methods). Thus, the exact number of personality factors and their interpretation will always be a somewhat arbitrary decision.

Personality Models Are Useful Fictions

Rather than seeing these "big few" personality models as unequivocal "truth" about personality structure, it is more appropriate to view them as useful abstractions. These models are valuable

tools for certain purposes, providing a simplified lens through which to explore personality variation.

By recognizing these models as useful tools, we can remain vigilant about their limitations. Relying solely on these broad abstractions can obscure more nuanced approaches to studying personality that may emerge in the future. Many scholars advocate for a shift toward investigating specific, fine-grained patterns of thoughts, feelings, and behaviors, recognizing that these may offer a more comprehensive understanding of individual differences.

For our purposes, however, these broad dimensional models will work well enough. They will provide us a common reference point as we learn about many topics in personality. The remainder of the course will explore these personality dimensions from many angles including evolution, genetics, physiology and neuroscience, motivation, development, and more.

References

Diener, E. & Lucas, R. E. (2023). Personality traits. In R. Biswas-Diener & E. Diener (Eds), Noba textbook series: Psychology. Champaign, IL: DEF publishers. Retrieved from <http://noba.to/96u8ecgw>

Feeney, M. G. (2020). Gordon Allport. The Wiley Encyclopedia of Personality and Individual Differences: Models and Theories, 1-9.

4 Measurement Matters

Some sort of quantification of the thing we want to study is necessary if we want to objectively compare things in a systematic way. Without measurement, we can't do science. So, after the Big Five was discovered, people began developing tools to specifically measure a person's standing on the traits of Extraversion, Agreeableness, Neuroticism, Conscientiousness, and Openness. Measurement of personality traits like the Big Five or HEXACO traits would, in theory, allow researchers to use their knowledge of a person's personality to predict things like behavior and life outcomes. The most common way of measuring personality traits in psychology is through lists of questions called "scales" or "inventories".

Personality inventories are typically made up of a set of questions or statements (items) that are designed to assess thoughts, behaviors, and feelings (e.g., "I often feel scared"). Respondents are typically asked to rate their agreement with each item on a Likert-style scale (e.g., strongly agree, agree, neutral, disagree, or strongly disagree); these categorical judgements are then assigned numbers by researchers for analysis (e.g., 1, 2, 3, 4, 5). The numerical values for responses to items that are intended to measure a common personality trait are aggregated together to obtain a score for each subject on the set of personality traits measured by the inventory.

But do these inventories measure the Big Five traits as intended?¹ How can we know? In this chapter, we'll discuss measurement matters—mainly reliability and validity—because measurement matters if we want accurate estimates of a person's personality. We'll first examine the concepts of reliability and validity. Then we'll explore some common tests for evaluating each and see how common Big Five inventories fare in these tests.

The Relationship Between Reliability and Validity

Before we get into the details of reliability and validity, let's examine how they relate to each other and to the process of measurement. Reliability is about being consistent – when we measure something more than once, we should get similar scores each time. Validity, on the other hand, ensures we're measuring what we actually want to measure. Both reliability and validity are key parts of accurate measurement.

¹The name for this chapter is shamelessly borrowed from: <https://www.psychologicalscience.org/observer/measurement-matters>

Reliability is like the first building block of measurement. We want our measuring tool to give us roughly the same estimate every time we use it, assuming the thing we're measuring hasn't changed. For example, we wouldn't want a ruler that gave different estimates of an object's length each time we measure it; and we don't want a personality inventory that gives drastically different estimates of a person's personality traits each time we measure them.

But getting reliable measurement just tells us that we can measure *something* consistently. It doesn't tell us what that something is. That's where validity comes in.

Validity helps us determine if we're really measuring what we intended to measure. For instance, if we're trying to measure Extraverted someone is, we want our measurement to reflect that trait and only that trait. We don't want our measurement tool to also estimate other traits like Openness or what our subject had for breakfast. Reliability is crucial, but it's not the whole story – we need validity too.

To put it simply, reliability is like the foundation of a good measurement tool. It sets the stage for consistency, but we also need validity to make sure we're measuring the right thing accurately. Put another way, reliability is *necessary but not sufficient* for good measurement. When we have both reliability and validity, however, we can be more confident that we've got ourselves a meaningful measurement tool and that the measurements we derive from it mean what we think they mean.

Reliability Assessment

How can we assess the extent to which a measurement tool gives consistent estimates? There are many methods for assessing reliability. We'll explore the main ones that have been used to examine the reliability of personality inventories: test-retest reliability, inter-item reliability, and inter-rater reliability.

All the reliability assessment methods we'll look at are based on correlations. I'm going to mostly assume here that you know what a correlation is, but here is a quick summary. Correlations summarize the overlap between variables. The correlation coefficient is bounded between -1 and +1. The closer a correlation is to 1, whether it's negative 1 or positive 1, the stronger the correlation is—or the more overlap there are between the two constructs. A correlation of zero means that there is absolutely no relationship or overlap between constructs, so the closer a correlation coefficient is to zero, the less strongly related they are. The positive and negative sign of the correlation coefficient tells us the direction of the relationship. If the sign is positive, then it means there is a positive correlation between variables, i.e., when one variable tends to be high, the other tends to be high as well. If the sign is negative, it means the correlation is negative, i.e., when one variable is high, the other tends to be low. When it comes to reliability assessments, we typically want to see strong correlations as evidence.

Test-Retest Reliability

Test-retest reliability assesses the stability and consistency of a measurement over time. It involves administering the same measurement to the same individuals on separate occasions. A strong positive correlation indicates that the measurement tool consistently captures something—but not necessarily the thing we want to measure—across different points in time. These measurement occasions can be separated by any amount of time deemed appropriate by the researcher, whether that’s minutes, hours, days, weeks, months, or years.

Assuming the researcher doesn’t think that personality changes much in the time frame, the scores obtained from the different measurement attempts should show high correlations if the measurement instrument is reliable. Note that assessing test-retest reliability depends on whether we think personality changes much and how quickly it might change.

For common Big Five inventories, the test-retest reliability is quite high. The table below shows the average test-retest reliabilities for each of the Big Five traits 74 studies (Gnambs, 2014). This kind of aggregation of results from many studies is called a meta-analysis.

Table 4.1: Test-retest reliability of Big Five measures

Dimension	Average test-retest r across 74 studies	SD of r across studies
Extraversion	0.834	0.081
Agreeableness	0.766	0.097
Conscientiousness	0.817	0.067
Neuroticism	0.802	0.089
Openness	0.802	0.077

Note: reproduced from Gnambs (2014).

Across all of the Big Five traits, the expected reliability of a test is about .82. That’s pretty high! Further, the small standard deviation (SD) values in the 3rd column indicate that there is not very much variation across studies. Many of these 74 studies used different inventories to measure the Big Five and had different time intervals between measures, but they all had generally high consistency. Tests that had more items tended to have somewhat better test-retest reliability.

This meta-analysis found that studies with longer times between tests typically had lower test-retest reliability than studies with less time between tests. But in general, reliabilities were still high, even for tests two months apart. In a later chapter we will explore more about how personality changes over time, so it isn’t necessarily a problem if reliability is lower over longer periods of time. We sort of expect that unless you think personality is perfectly fixed across time (some researchers do).

Inter-Item Reliability

Inter-item reliability evaluates the internal consistency of a multi-item scale or questionnaire designed to measure a specific construct. It examines how well the individual items within the scale correlate with one another. This is typically assessed using methods like Cronbach's alpha, which you'll see in many papers indicated by its Greek letter, α . A high inter-item correlation suggests that the items are repeatedly measuring a similar thing—ideally the same thing. If we have very low correlations among our items, it may suggest that they are not good items.

Note that inter-item reliability is largely dependent on the number of items we put in a scale and how similarly they are worded: scales with more items, especially more similarly phrased items, will tend to have stronger inter-item reliability. This means that if we are trying to measure several partially distinct aspects of a personality trait with only a few items, our inter-item reliability might be low even though the questions do a good job of measuring the trait. Even though higher inter-rater reliability is typically better, we need to think carefully about what we are intending to measure if we want to make good judgments about inter-item reliabilities.

For most Big Five measures, inter-item reliability is very high. The table below shows the inter-item reliability coefficient (alpha, α) for the IPIP-50, which I administered to students on the first day of class. This high reliability is typical of most Big Five scales, especially longer ones. Similar alpha coefficients are found for most HEXACO inventories, except for ones with very few items per dimension (e.g., de Vries, 2013).

Table 4.2: Inter-item reliability for the in-class Big Five personality test.

Dimension	Inter-item reliability (α)
Extraversion	0.91
Agreeableness	0.70
Conscientiousness	0.78
Neuroticism	0.77
Openness	0.80

Inter-Rater Reliability

Inter-rater reliability gauges the level of agreement or consistency among different raters or observers when evaluating a particular phenomenon or behavior. People can have different perceptions of most psychological outcomes, especially personality, so we want to be able to assess how much estimates from personality inventories differ from person to person. High inter-rater reliability scores signify a high level of agreement among the raters, enhancing the confidence in the accuracy and consistency of the measurements.

When it comes to personality inventories, most inter-rater reliability estimates come from having two (or more) people rate one other person. Or comparison's will be made between a person's self-assessments on the inventory with assessments of that person made by others who know them (e.g., friends, coworkers, romantic partners).

For the Big Five personality inventories, inter-rater reliability is typically moderately high. Across a range of studies that compared the consistency between different peoples' perceptions of another person's personality, the average inter-rater reliability was about .4 (Vazire & Carlson, 2013). That means that inter-rater reliability is far from perfect (a correlation of 1). However, it makes sense that people would have different assessments of another person's personality—we all have different bits of incomplete information that we interpret using our own perspectives. What's more reassuring is that relationship length and closeness seems to predict stronger agreement between a person's self-rated personality and their friend's or romantic partner's ratings (Watson et al., 2000).

Interestingly, even in a first meeting, people may even be able to judge others' personality with some accuracy. To illustrate this, I have students talk to each other a little bit on the first day of class without any real guidance or direction. After talking, each pair of students rate each other on each of the Big Five dimensions. I then compare those ratings from brief interactions to the students' self-ratings on a longer personality test. The table below shows the correlations between students' self-rated personality traits and their partners' perceptions of their personality from their brief interactions. Clearly, the accuracy of these students' perceptions is not very high, but also it is indeed better than chance in most cases (e.g., extraversion and agreeableness).

Table 4.3: Correlations among self-rated personality traits and peer-rated personality traits from a brief interaction

	Self-rated Extraversion	Self-rated Agreeable- ness	Self-rated Conscientious- ness	Self-rated Neuroticism	Self-rated Openness
Peer-rated Extraversion	0.47	0.35	0.23	0.01	0.33
Peer-rated Agreeableness	0.15	0.26	0.19	0.15	0.07
Peer-rated Conscientious- ness	0.22	0.17	0.00	0.18	0.12
Peer-rated Neuroticism	0.11	0.16	0.06	0.24	0.01
Peer-rated Openness	0.22	0.25	0.15	0.18s	0.08

Note: Based on a sample of 112 students from three classes across two semesters.

There is a whole huge area of research exploring how and why people's perceptions of others' personality may differ and how these perceptions change over time, which we may be able to explore in a later chapter. For now, we should just recognize that self-other and even other-other agreement is not very high, but that doesn't necessarily mean there is something wrong with the measures themselves.

Validity Assessments

So, the reliability of most Big Five tests that are commonly used is typically pretty good. There is strong test-retest reliability across fairly long time periods, there is strong inter-item reliability, and there is moderate inter-rater agreement. I'd be comfortable saying that most Big Five tests are reliable, meaning they can consistently measure something. But again, reliability is not enough—we want to know that we are consistently measuring the specific something that we intend to measure (i.e., Big Five traits). Let's look at some methods for assessing validity and see what the evidence for the validity of Big Five personality inventories looks like.

There are two broad types of validity assessments: *subjective* and *empirical*. Subjective validity tests include face validity and content validity. Empirical validity tests include predictive validity, concurrent validity, and discriminant validity. Figure 4.1 depicts the relationship between specific validity assessments and each type of validity. We'll explore each of these in more detail in their respective sections below.




Figure 4.1: Different types of validity.

Subjective Validity Assessments

Subjective validity can be assessed without collecting any data or doing any statistical tests. We can all conduct a subjective validity test for any inventory just by looking at it and thinking a bit. We just need a good sense of the *operational definition* of the trait. The operational

definition describes the specific thing being measured (e.g., Extraversion) in a way that allows us to determine how the measurement tool can achieve measurement. Let's take a closer look at face validity and content validity.

Face Validity

Face validity is an initial, subjective evaluation of whether a measurement tool appears, “on its face”, to measure the intended construct accurately. It assesses whether the items or questions in the instrument seem relevant and appropriate to the construct being measured.

Assessing face validity is like judging a book by its cover. It's a quick, intuitive assessment of whether a measurement tool appears to measure what it's supposed to. Just as you might get a sense of a book's content by looking at its cover and title, face validity is about the inventory items looking like it measures what it is supposed to measure.

However, face validity isn't always desirable. Sometimes we don't want subjects to know exactly what is being measured. For example, if we were to try to measure psychopathic traits using a self-report scale, we might not want it to be obvious because then a good psychopath could just lie to us if they don't want to be measured. However, we should always consider the face validity of our measure because it is likely that our subjects will make some guesses about what we are measuring and those guesses can affect their responses—so by considering what it looks like we are measuring, we might be able to better interpret participant responses.

Content Validity

Content validity is a systematic and comprehensive assessment of how well the items in a measurement tool represent the construct being measured. It involves judgment to ensure that the items adequately cover all aspects of the construct.

Content validity is like making sure you have all the pieces of a puzzle that you are trying to put together. If you are missing jigsaw pieces, the puzzle will, at best, be incomplete; and if you're missing too many pieces you may not even be able to produce anything that looks like the picture you intended to produce (i.e., measure the thing you intended to measure). We want to make sure our questions in an inventory (corresponding to puzzle pieces) provide complete coverage of our trait (corresponding to the puzzle picture).

Empirical Validity Assessments

Empirical tests require us to collect some data and run some statistical tests. This doesn't necessarily mean they are purely objective tests; researchers must still make subjective interpretations of the statistical results to determine what sort of evidential value they have.

Let's take a closer look at the empirical validity assessments: criterion/predictive validity; concurrent/convergent validity; and discriminant/divergent validity.

Criterion (Predictive) Validity

Criterion validity, also known as predictive validity, refers to the extent to which a measurement tool accurately predicts or correlates with a specific outcome that theory predicts it should be closely related to. It involves comparing the scores obtained from the measurement tool with the scores on an established criterion, and ideally demonstrating a strong relationship between them.

Criterion validity can be likened to hitting the bullseye in darts or archery. Imagine hitting the center of the target accurately is your goal. Criterion validity provides a way to measure how close or far your throws (measurements) come to hitting the bullseye. Or if you have a crystal ball that is supposed to predict the future, the predictive validity would be indexed by how many times the crystal ball actually predicts the future.

For personality inventories, the predictive utility typically comes from looking at how well scores on inventories predict behavior that should be relevant to a given trait. There have been many studies looking at this.

For example, in 1989 David Buss (my PhD advisor) and Mike Botwin (my colleague at Fresno State), asked people to report how often they engaged in various behaviors that were theoretically reflective of different personality traits. Then they compared the behavioral reports to the scores on a Big Five inventory. They found that the behaviors that people reported engaging in tended to be moderately correlated with relevant inventory scores for the Big Five traits. The table below gives some examples of the behaviors that go with each trait and their specific correlations with each.

Table 4.4: Correlations between Big Five scale scores and reported behaviors

Big Five Scale	Example Criterion Acts	Mean r
Extraversion	Dancing in a crowd, telling jokes	0.51
Agreeableness	Offered to help a friend move, glared at a stranger (-)	0.53
Conscientiousness	Didn't review work (-), saved money	0.54
Neuroticism	Worried about something out of control, put self down	0.37
Openness	Went to an art exhibit, discussed an issue from multiple points of view	0.45

Note: reproduced from Buss & Botwin (1989)

Other studies have looked at even more stringent behavioral criterion by measuring behavior directly, rather than relying on self-reports. This is important because it is possible that self-reported behaviors are biased by personality. For example, extraverts could be more likely to remember and report talking to people. This may lead us to overestimate correlations between behavior and personality scores.

Indeed, when we examine studies that directly measure behavior, the correlations with inventory scores are typically smaller. A meta-analysis of correlations between personality test scores and directly measured behavior—either in lab settings or in more naturalistic settings—found that the average correlation between an inventory score for each of the Big Five and an associated behavior was around .25 (Vazire and Carlson, 2010). This is quite a bit smaller than the self-reported behavior correlations, but it does still provide evidence that self-reported inventory scores are picking up on actual differences in behavioral tendencies.

Concurrent (Convergent) Validity

Concurrent validity—sometimes called convergent validity—involves comparing a new inventory with an existing, well-established inventory that measures the same construct. To do this comparison, the different inventories are typically administered to a group of people at the same time, so that we have scores for everyone on both inventories, allowing us to assess their correlation. Ideally, we want very high correlations on inventories that are supposed to measure the same thing. For example, different Neuroticism inventories should be highly correlated, as should different Openness measures.

Maybe you can find it easier to remember concurrent validity by thinking about its root word “concur”. If someone—like an old-timey judge in a Mississippi courtroom—says “I concur” (in a thick southern accent), it means they agree. When we examine concurrent validity, we are looking for measures of the same construct to agree with each other in the estimates they produce. Alternatively, for remembering convergent validity, you can use the root “converge”. When rivers come together, we say they converge. When we assess convergent validity, we want our two measures to come together on the same estimates.

When it comes to personality inventories, convergent validity is typically assessed by comparing new measures to some older established Big Five inventories. For example, when the HEXACO personality inventory was created, the authors had to establish that their inventories for traits that had clear overlap with the Big Five traits (e.g., Extraversion, Openness), were indeed highly correlated with those traits. So, they administered their HEXACO inventory to people along with an existing Big Five inventory and found the correlations shown in the table below.

Table 4.5: Correlation between HEXACO inventory scores (rows) and Big Five inventory scores (columns)

HEXACO scales	Neuroticism	Extraversion	Agreeableness	Conscientiousness	Openness
Honesty-Humility	-0.08	-0.14	0.28	0.09	0.13
Emotionality	0.55	0.01	0.34	-0.06	0.11
Extraversion	-0.14	0.74	0.13	0.1	0.21
Agreeableness	-0.37	-0.16	0.52	-0.09	0.03
Conscientiousness	0.02	0.03	0.01	0.7	0.2
Openness	-0.03	0.07	0.12	-0.1	0.76

Note: reproduced from Lee & Ashton (2004)

The most relevant correlations to examine in this correlation matrix are the bolded correlations that depict the correlations between each HEXACO trait and its Big Five counterpart (e.g., HEXACO emotionality and Big Five Neuroticism). For the most part, there are strong correlations between the HEXACO traits and their corresponding Big Five traits.

Another application of convergent validity is when researchers make a shorter scale to measure Big Five traits. Older Big Five inventories typically have many questions—sometimes hundreds—and they take a long time for participants to fill out. That makes it hard to get a lot of good data because most people don’t want to spend a long time filling out surveys. So, researchers often try to make shorter scales to reduce the burden on participants and make it easier to collect data. But when we do that, we need to understand how it changes our measurements. In order to do that, researchers will typically administer the new shortened inventory in tandem with a longer inventory to assess the correlations among their scores. For example, Gosling et al. made a very short 5-item scale that measures each Big Five trait with just one item. They found that the scores from single items were moderately correlated with scores from a longer inventory (ranging from around .8 for Extraversion to .48 for Openness). This information can help researchers evaluate the tradeoffs of longer versus shorter measures in their own research.

Discriminant (Divergent) Validity

Discriminant validity—sometimes called divergent—is the degree to which a measurement tool accurately distinguishes between the construct it’s intended to measure and other unrelated constructs that it isn’t supposed to measure. It ensures that the tool does not measure other

irrelevant traits, demonstrating its ability to discriminate between different concepts or variables.

You may find it helpful to remember that discriminant validity has the root “discriminate”. While it generally isn’t cool to discriminate (e.g., against people), when it comes to measurement, we definitely want to discriminate. That is, want our measurement inventories to discriminate against the traits that we are not trying to measure by not measuring them—only the trait we want to measure should be allowed into our estimates of the trait.

For personality inventories, discriminant validity is often used to argue that a given inventory measures something different from what is already captured in the Big Five. For example, the HEXACO creators argued for the importance and novelty of the Honesty-Humility dimension by showing that it was largely uncorrelated with scores on Big Five dimensions (see the correlations in Table 4.5 above). This suggests that it is capturing something new. Additionally, we can learn from HEXACO correlations with Big Five scores that Agreeableness and Emotionality in the HEXACO inventory are somewhat different from Agreeableness and Neuroticism in the Big Five inventory, since they share only about half of their variance.

It is also important that Big Five scores are mostly uncorrelated with one another as well—we don’t really want Extraversion scores to be picking up on Neuroticism or Openness because that makes it harder to interpret what they mean. In reality, however, it is often the case that scores in different personality dimensions are correlated at least a little bit, perhaps due to some self-report biases. This phenomenon is often attributed to a mythical “crud” factor that accounts for confounds that might drive correlations among scores from inventories in psychology.

What’s the Verdict? Do Personality Inventories Measure What They Intend To?

In this chapter we’ve examined the important components of measurement, reliability and validity. We learned how to assess information about inter-rater reliability, test-retest reliability, and inter-item reliability, which is necessary but not sufficient for measurement. Then, we learned how to evaluate the validity of a measure by testing criterion validity, convergent validity, and discriminant validity.

I gave you some representative evidence for each of these aspects of reliability and validity for Big Five and HEXACO inventories. At this point, you should be able to form your own informed opinion about what these personality inventories measure, as well as evaluate future personality measurement tools using the same criteria. But I’d encourage you to avoid being too dichotomous about our decisions. Just like personality traits themselves, reliability and validity exist on a continuum and we have to consider each measurement tool in relation to others that are available.

In my opinion, there seems to be pretty strong reliability indicating that the inventories measure something fairly consistently. And there is some evidence that scores on these inventories predict real world behaviors that are relevant to each personality trait, albeit far from perfectly. A skeptical but reasonable interpretation could be that these inventories provide useful information about personality, but they don't necessarily measure personality directly. That is, they are useful proxies of personality traits.

I think we should view existing personality inventories as practically useful tools, but we should not stop trying to develop more precise measurement methods. It seems unlikely to me that the common inventory approach—which was pretty much our first attempt at systematically measuring personality traits—will be the best we can do. We should continue to explore potential alternatives. Maybe you can help us develop some!

But for now, most of what we know about personality is based on inventory-like measures. So, most of the research we'll talk about throughout the rest of the course will largely be based on these sorts of measures. We should keep this in mind as we evaluate the findings from different areas of research.

References

Gnambs, T. (2014). A meta-analysis of dependability coefficients (test–retest reliabilities) for measures of the Big Five. *Journal of Research in Personality*, 52, 20-28.

Lee, K., & Ashton, M. C. (2004). Psychometric properties of the HEXACO personality inventory. *Multivariate behavioral research*, 39(2), 329-358.

Vazire, S., & Carlson, E. N. (2010). Self-knowledge of personality: Do people know themselves?. *Social and personality psychology compass*, 4(8), 605-620.

Watson, D., Hubbard, B., & Wiese, D. (2000). Self–other agreement in personality and affectivity: The role of acquaintanceship, trait visibility, and assumed similarity. *Journal of personality and social psychology*, 78(3), 546.

5 Personality Power

Thus far, we've explored how personality taxonomies like the Big Five or HEXACO came to be and critically examined our ability to measure relative differences between people on personality dimensions. We learned that contemporary personality taxonomies are largely informed by the natural-language adjective terms that people use to describe others personality. We also scrutinized the measurement tools used to measure personality traits. We found decent evidence of their reliability and validity, suggesting that personality inventories can be useful proxies for differences in personality.

But just because we can measure something, doesn't mean it's important. Given that personality traits are theoretically expected to be stable characteristics of a person that influence how they interact with the world, we might expect that personality traits could allow us to predict and explain many aspects of a person's life. So, is personality powerful? In this chapter, we'll discuss the evidence for the power of personality in prediction and explanation.

The Person-Situation Debate

We can't really talk about the power of personality without treading into an age-old debate. People have been arguing for ages about whether internal or external forces are more important for understanding human behavior. In contemporary psychology, this dispute manifested in what has come to be known as the person-situation debate. There are two opposing sides to this debate.

Person (Trait) Position

Proponents of the person side of the debate—often referred to as trait theorists—argue that individuals possess consistent and enduring traits that influence their behavior across situations. These traits, like extraversion or conscientiousness, are seen as intrinsic to the person and relatively stable over time. This perspective emphasizes the importance of understanding an individual's unique characteristics to predict their actions.

As would be expected from this position, proponents of the person side predict that people should behave relatively consistently across situations and across time. The main research focus on this side of the debate is to identify and study personality traits.

Situation Position

Advocates of the situation perspective emphasize the powerful influence of external circumstances and contexts on behavior. They argue that people respond differently based on the situation they are in, regardless of their personality traits. This perspective emphasizes the power of the situation for shaping human behavior.

Following from this position, situationists would expect that people's behavior is not consistent across time and across situations. Accordingly, the main research focus on this side of the debate is to identify situations and features of situations that influence human behavior.

Walter Mischel's Critique of Personality

At the height of the person-situation debate, Walter Mischel challenged the traditional view of personality by questioning the stability and consistency of behavior across different situations. He argued that the assumption of consistent, enduring traits guiding behavior oversimplifies the complexity of human actions. Mischel's critique was encapsulated in his book "Personality and Assessment" (1968).

Mischel emphasized the importance of understanding behavior in specific situations, arguing that individuals do not possess universal traits that dictate their actions across diverse contexts. He conducted extensive research showing that people often display different behaviors in different situations, indicating that behavior is highly dependent on the specific circumstances.

Furthermore, Mischel highlighted the role of cognitive processes, such as perceptions, beliefs, and interpretations, in shaping behavior. He argued that an individual's interpretation of a situation and their expectations of outcomes significantly influence how they behave in that particular context. These cognitive processes were seen as essential components in understanding the variability in behavior.

Mischel's "person-situation interaction" theory proposed that behavior is an outcome of the interplay between a person's individual characteristics (such as cognitive and affective processes) and the situational cues. He emphasized that individuals possess a range of behavioral potentials, and their behavior is activated based on the unique demands of the situation.

Overall, Mischel's arguments challenged the conventional view of personality as a set of stable traits, paving the way for a more nuanced and contextually sensitive understanding of human behavior. His work encouraged researchers to explore the intricate relationship between individual differences and the environments in which they manifest, fostering a more comprehensive and dynamic perspective on personality psychology.

The Power of the Situation

Following the second world war, social psychologists conducted a lot of studies that highlighted the power of the situation for shaping how people behave. The most famous of these studies are probably the Milgram Experiments and the Stanford Prison Experiment.

The Milgram experiment was conducted by Stanley Milgram in the early 1960s to investigate obedience to authority figures. Participants were instructed to deliver electric shocks to a learner in another room when they answered questions incorrectly, with the shocks increasing in intensity each round. The shocks were actually fake and the learner was actually an actor in on the experiment, but the participants didn't know that. The experiments showed that around 60% of participants obeyed authority and administered seemingly harmful shocks, even when they felt distressed by their role. This demonstrated the power of authority and social pressure in influencing individuals to act against their own moral inclinations.

The Stanford Prison Study was conducted by Philip Zimbardo in 1971 to investigate the psychological effects of perceived power and authority in a simulated prison environment. Participants were randomly assigned to be either guards or prisoners in a makeshift prison in the basement of the psychology department building. The study was supposed to last two weeks but was halted after only six days due to the extreme psychological distress and abusive behavior exhibited by the guards towards the prisoners. The results are often interpreted as evidence that social roles and the institutional environment can profoundly influence individuals' behavior.

When they aren't providing fodder for many a discussion on ethics in psychology research, these studies are typically invoked to highlight the power that the situation has on people's behavior and to downplay the importance of personality. But do they really discount the importance of personality traits?

Follow up research has examined how personality traits may have played a role in both of these famous studies. For example, researchers using a modified version of the Milgram Experiment measured participants' personality traits and found that participants with higher Conscientiousness and Agreeableness scores were more likely to adhere to the experimenter demands (Bègue et al., 2015). And other researchers have found that people who would voluntarily sign up for a study on "prison life", like the Stanford Prison study, tend to be higher on aggressiveness, authoritarianism, Machiavellianism, narcissism, and social dominance—and lower on empathy and altruism—than would be expected if there were no bias in who signs up for such a study (Carnahan & McFarland, 2007).

Together, these follow up studies suggest that the personality traits of participants likely played some role in the results of the original studies. While the Milgram Experiments and Stanford Prison Study may highlight the impact that extreme situations can have on behavior, they certainly do not rule out a role for personality.

The Power of Personality

The growing evidence for the power situations spurred research examining the predictive and explanatory power of personality traits. We now have evidence that the personality estimates we can get by administering personality inventories can be useful for predicting important outcomes.

For example, Ozer and Benet-Martinez (2006) conducted a qualitative assessment of associations between personality and consequential life outcomes. They found that each of the Big Five personality traits could be linked to a range of outcomes that people care about including well-being, relationship satisfaction, drug use, and religiosity.

Subsequently, Roberts et al. (2007) conducted a meta-analysis to examine the correlation between personality traits and specific outcomes in comparison to other popular predictors. They found that personality traits were better predictors of mortality risk and divorce risk than intelligence test scores and socioeconomic status. When predicting job performance, only intelligence test scores were a better predictor than personality traits. So, when it comes to relative predictive power, personality traits may be some of the best predictors we have.

More recently, Beck & Jackson (2022) used data on over 170,000 people from 10 different longitudinal studies to test the power of personality. They conducted mega-analyses that controlled for many different sets of background variables (e.g., race, age, socioeconomic status). They found that personality trait scores were correlated with 14 life outcomes (e.g., divorce, marriage, having children, losing a job) and that controlling for potential confounds or moderating characteristics did not substantially dampen their predictive power. You can explore all the analyses for each personality trait and life outcome on the website they created to display their results: <https://emoriebeck.shinyapps.io/Big-5-Prediction/>.

Aside: Prediction versus (Causal) Explanation

These studies that highlight the power of personality are based on cross-sectional observational correlational data, which makes the causal picture unclear. We should take a step back and remember the adage: correlation is not causation. It is easy to interpret correlations between personality and life outcomes as indicating that personality traits cause life outcomes. But this would be a leap of logic that is not warranted by the data and evidence we currently have. Something can be a good predictor of an outcome and not actually have any causal effect on the outcome. Future research may be able to use recent advances in causal inference to better test the causal effects of personality on life outcomes. But for now, we can still appreciate the predictive power of personality.

Resolving the Person Situation Debate

After years of heated debate, we can now declare the person-situation debate over—or at least not worth discussing further (Hogan, 2009). The person-situation debate is largely a waste of time because neither extreme is likely to be true. Most contemporary psychologists agree that behavior in any given instance is an interaction between features of the situation and features of the person.

In 2008, personality researchers William Fleeson and Erik Nofle provided a way to synthesize both the person and situation to end the person-situation debate. They showed that peoples' behavior does indeed vary substantially across situations, but that people exhibit stable individual differences in the unique distributions of behaviors across time and contexts. Thus, we can understand personality as the person-specific distribution of behaviors across contexts. They further showed that the characteristics of these distributions, such as the mean or variance, in different periods of time are highly correlated when broken up. For example, if we measure peoples' average extraversion levels across situations for one week, we will see that they are strongly correlated with average extraversion levels across situations for the next week. The longer the observation period that is used to make the average, the stronger the correlations are with future intervals of the same length. This research highlights that behavior can change drastically from situation to situation, but there can still be a signature of stable personality traits.

Our understanding of human psychology is far from complete. We must recognize the importance of both situations and personality traits—and their interactions—to predict and explain human behavior. The future of personality research likely involves specifically measuring aspects of situations and aspects of people to examine how and why people differ in their characteristic reactions to situations (e.g., Kuper et al., 2022).

References

- Beck, E. D., & Jackson, J. J. (2022). A mega-analysis of personality prediction: Robustness and boundary conditions. *Journal of Personality and Social Psychology*, 122(3), 523.
- Bègue, L., Beauvois, J. L., Courbet, D., Oberlé, D., Lepage, J., & Duke, A. A. (2015). Personality Predicts Obedience in a Milgram Paradigm. *Journal of personality*, 83(3), 299-306.
- Carnahan, T., & McFarland, S. (2007). Revisiting the Stanford Prison Experiment: Could participant self-selection have led to the cruelty?. *Personality and Social Psychology Bulletin*, 33(5), 603-614.

- Fleeson, W., & Nofhle, E. (2008). The end of the person–situation debate: An emerging synthesis in the answer to the consistency question. *Social and Personality Psychology Compass*, 2(4), 1667-1684.
- Hogan, R. (2009). Much ado about nothing: The person–situation debate. *Journal of Research in Personality*, 43(2), 249.
- Kuper, N., Breil, S. M., Horstmann, K. T., Roemer, L., Lischetzke, T., Sherman, R. A., ... & Rauthmann, J. F. (2022). Individual differences in contingencies between situation characteristics and personality states. *Journal of Personality and Social Psychology*, 123(5), 1166.
- Ozer, D. J., & Benet-Martinez, V. (2006). Personality and the prediction of consequential outcomes. *Annu. Rev. Psychol.*, 57, 401-421.
- Roberts, B. W., Kuncel, N. R., Shiner, R., Caspi, A., & Goldberg, L. R. (2007). The power of personality: The comparative validity of personality traits, socioeconomic status, and cognitive ability for predicting important life outcomes. *Perspectives on Psychological science*, 2(4), 313-345.

6 Evolution and Personality

Now that we have a decent taxonomy of personality traits for focusing our attention (i.e., the Big Five), we can begin to ask interesting questions about them. In this chapter we will turn our focus to an *ultimate* question: Why do these personality differences exist?

To begin to answer this question, we need to first have an understanding evolution. This is because evolution is the best scientific theory developed for explaining why living things are the way they are. Humans have long observed that living organisms, like plants and animals, are typically very well suited to survive and reproduce in the habits they live in. It often seems as though organisms are *designed* for their specific habitats. Evolutionary theory provides us a framework for understanding why this is the case.

Darwin's Big Ideas

The theory of evolution was developed by Charles Darwin in the mid 1800s. He spent years traveling around the world and observing the wondrous diversity of animals and plants across many exotic habitats, like the Galapagos Islands. Based on his observations, he argued that all living organisms share a common ancestry and proposed that a simple process could explain the vast diversity of organisms we observe today. The simple process he proposed is called *natural selection*.

Natural Selection

Natural selection depends on three conditions: *variation*, *differential reproduction*, and *inheritance*. In self-replicating populations—whether that self-replication is sexual or asexual—these three ingredients inevitably lead to changes in the frequencies of traits within the population over time. We'll explore what these conditions are in detail below.

Variation

In any population (e.g., a crash of rhinos, an exaltation of larks, a tower of giraffes¹), individuals are slightly different from one another. Individuals within populations differ from one another

¹Yes, these really are the names for groups of these animals!

in countless ways: body size, musculature, hunger, adventurousness, aggressiveness, furriness, speed, sleepiness, etc. All the traits that an organism possesses make up its *phenotype*. So, individuals within populations have somewhat different phenotypes. This phenotypic variation eventually leads to some traits becoming more prominent in a population because of differential selection and inheritance.

Differential Reproduction

Some individuals within populations reproduce more than others and some don't reproduce at all. Sometimes individuals fail to reproduce because they die from some completely random accident. But individuals also fail to reproduce because there is something about their phenotype that led them to die before reproduction or to be unsuccessful in finding or copulating with a mate. Put another way, some individuals possess traits that make it harder to survive and reproduce in their habitat, while some individuals have traits that make it easier to survive and reproduce. For example, a slower adolescent gazelle might be more likely to be caught by a cheetah before they have a chance to mate; or a more brilliantly colored bird might be more noticeable than a duller bird, leading them to be detected more easily by potential mates and giving them more reproductive opportunities than duller birds. The degree to which an individual organism has traits that allow it to successfully survive and reproduce in their environment better than other individuals is referred to as differences in *fitness*.

Some of the phenotypic variation in the population is causally linked to survival and reproduction. In other words, some individuals possess phenotypes that make it more likely that they will survive and reproduce, while others possess phenotypes that make it less likely that they will survive and reproduce. This is where the "selection" in natural selection takes place. But variation and differential reproduction alone isn't enough to produce evolution without the traits being passed down to the next generation of replicates in some way.

Inheritance

The final condition necessary for evolution to happen is inheritance. Inheritance allows the phenotypic traits that produce better chances of survival and reproduction (i.e., the traits with more fitness) to be passed on to the next generation of replicates (e.g., offspring). That means the traits must be *inheritable*.

Humans have recognized for a very long time that the offspring tend to resemble their parents. So, we've known for a long time that parents could pass on some traits to offspring. In Darwin's time, however, they didn't know about genes or even how exactly traits got passed on—in fact they thought that parents' traits got blended somehow, like mixing paints. But Darwin knew this wouldn't be enough for the fitness-enhancing traits to be preserved generation after generation. It turns out that around the same time that Darwin was developing his theory, Gregory Mendel was actively conducting research on inherited traits in pea plants which

ultimately led to our modern understanding on how traits are passed on generation after generation.

Nowadays, we know that parents' traits can be inherited by offspring via *genes*. Due to a process called recombination, offspring receive a random set of genes from each parent, which allows traits to be preserved across generations under certain conditions. Thus, parents who possess traits that make it easier to survive and reproduce will tend to produce offspring that possess those traits as well.

Natural Selection in Action

When the three conditions of variability, differential reproduction, and inheritance are present in a population of self-replicating individuals, fitness-enhancing traits tend to become more common. Usually, the changes in a population occur very slowly across many generations. Sometimes they can happen faster.

A famous example evolution is finches on the Galapagos Islands that inspired Darwin's theory. These finches exhibit remarkable variation in their beak shapes, and this variation corresponded closely to their dietary preferences and feeding behaviors. Some had slender, probing beaks suitable for capturing insects, while others had sturdy, crushing beaks for consuming seeds or cracking nuts. These variations in beak shape were advantageous for the finches' survival and reproduction within their ecological niche.

Over generations, this process of natural selection led to an increase in the frequency of the advantageous beak shapes within each finch population. These changes occurred relatively slowly but were nonetheless observable over time. Darwin's observations of the finches and their adaptive evolution provided crucial evidence for his theory of natural selection, demonstrating how small, gradual changes in traits within a population could lead to significant evolutionary shifts over extended periods.

Sexual Selection

Darwin's was a careful thinker and was always looking for ways to disprove his theory evolution. Apparently, he was somewhat tormented by the existence of traits that do not appear to directly support survival or even hinder survival, such as the peacocks long and cumbersome tailfeathers. How could something like that evolve if natural selection was the only process operating on traits?

Eventually, Darwin came up with a solution to this puzzle. He proposed a second evolutionary process, called *sexual selection*, which operates somewhat differently than natural selection. Sexual selection is a specialized form of natural selection whereby characteristics that improve success in competing to attract a mate become more common in a population, even if those traits aren't directly useful for survival—sometimes, traits can even harm survival and still be

maintained via sexual selection. There are two mechanisms of sexual selection: *intra*-sexual competition and *inter*-sexual choice.

Intra-Sexual Competition

Intrasexual competition involves members of one sex (usually males) competing for access to the opposite sex (usually females). This competition can take various forms, such as physical combat, displays of strength, or contests of endurance. The winners of these competitions often gain preferential access to mates and reproductive opportunities.

Because of this competition being directly tied to reproductive success, any traits that provide a competitive advantage will tend to become more common in a population, assuming that the traits are inheritable. So, over long periods of time, traits like antlers, horns, and thicker skin have evolved in the sex that typically competes for mates (usually males) across many species. Many of these traits do not directly aid in survival. For example, male deer only grow antlers during the mating season, indicating that they are not used for protection in a general sense.

Inter-Sexual Choice

Intersexual choice, also known as mate choice, involves members of one sex (often females) choosing mates based on certain preferred characteristics. These characteristics, often referred to as “ornaments” or “indicators of fitness,” can include physical traits, behaviors, or social status. Sometimes, however, the preferences are somewhat arbitrary. Individuals with preferred traits are more likely to be chosen as mates, allowing them to pass their traits to the next generation.

Sexual selection often leads to the development of exaggerated traits that are not directly related to survival but are favored because they appeal to the opposite sex. This can result in elaborate courtship displays, colorful plumage, or intricate behaviors that enhance an individual’s attractiveness to potential mates. The peacock tail and complex bird songs are classic examples of traits that likely evolved via inter-sexual choice.

Products of Evolution

Evolution is ultimately just a change in frequencies of traits (and their associated genes) in a population over time. There are three main products of evolutionary processes: adaptations, byproducts, and noise.

Adaptations

Adaptations are traits that evolve because they aided in survival and reproduction. Put another way, adaptations are functional, like a tool. Adaptations are for doing things. Because traits that reliably aid in survival and reproduction tend to become more common in the population, adaptations are generally species-typical and reliably developing—all or most members of a population (e.g., species) will develop the trait. An example of an adaptation is the umbilical cord that transports nutrients to a fetus from its mother during gestation.

Byproducts

Byproducts are traits that are not directly shaped by evolutionary processes, but they are somehow linked to adaptations. These may have initially emerged due to selection for a different trait or due to genetic or developmental constraints. Byproducts can persist in populations as they do not impose a significant fitness cost and are not actively selected against. Byproducts are non-functional, meaning that they do not directly aid in survival and reproduction. An example of a byproduct is the bellybutton. The bellybutton doesn't serve any function itself after gestation. It is a byproduct of the umbilical cord.

Noise

Noise refers to variation in traits that evolves and is not selected out by evolutionary processes because it is not tied to survival and reproduction. This noise is not part of the design of an organism. For example, peoples' bellybuttons vary in terms of size and appearance, but the size and appearance of bellybuttons doesn't matter for survival and reproduction on average. Noise, or random variation in traits, is inevitable because of random genetic mutations and environmental influences on traits. This variation is also a key driver of evolution because it can produce variants of traits that improve fitness (i.e., survival and reproduction).

The Power of Evolutionary Theory

Understanding the mechanisms and products of evolution, particularly natural and sexual selection, unveils the intricate tapestry of life. It allows us to comprehend how and why living organisms seem to fit so well within their natural habitats. Natural selection illustrates how subtle variations, advantageous or otherwise, influence survival and reproduction, shaping species over generations. Moreover, sexual selection illuminates the role of mate preferences and competition, driving the development of striking traits and behaviors that might seem unrelated to survival but profoundly impact a species' evolutionary trajectory. Together, these mechanisms produce adaptations, byproducts, and noise. Evolutionary theory provides a powerful lens through which we can fully appreciate the breathtaking diversity and complex

interplay of life on Earth. In the next section, we will explore how it may also allow us to better human psychology as well.

Applying Darwin’s Ideas to Humans: Evolutionary Psychology

Evolutionary Psychology (EP) is relatively new approach to studying the mind and behavior. At its core, EP is about applying evolutionary theory to understand the mind and behavior. This approach posits that many—but not necessarily all—aspects of human psychology and behavior exist because they enhanced our survival and reproduction in ancestral environments. The main focus of EP are *evolved psychological mechanisms*.

What the #@%! Is a Psychological Mechanism?

A psychological mechanism is an abstract term to describe a hypothetical information processing system in the mind. I like to think of psychological mechanisms as biologically based computer programs that are each designed to do a particular task. Psychological mechanisms are typically hypothesized to take some inputs and produce some outputs (e.g., behavior) based on some set of (logical or numerical) algorithmic procedures. So, they have a structure like the one depicted in Figure 6.1.




Figure 6.1: Diagram of a psychological mechanism.

The **inputs** to a psychological mechanism could be features of a situation, sounds, light refracting off objects. It depends on what the psychological mechanism evolved to do (i.e., what its function is). The inputs could even be outputs from other psychological mechanisms.

The psychological mechanism then performs some **computation** on the input(s). Just like a computer program, the computation depends on what the mechanism is designed to do. Some algorithms might be analogous to adding numbers together; some might be more like conducting a regression analysis; and some might be more logical rules like “if snake is detected, run away”.

Ultimately, these computations produce some **outputs** in the form of behaviors, thoughts, or information that is transmitted to other psychological mechanisms. Evolutionary psychology posits that evolved psychological mechanisms are specially designed to produce outputs that would have helped our ancestors survive and reproduce. Put another way, the outputs are supposed to help an organism solve an adaptive problem—we'll unpack this term in the next section.

To fully understand a psychological mechanism, we would want to eventually be able to provide a description for each of Marr's levels, which we covered in the Foundational Frameworks chapter. The computational level explanation would describe what the psychological mechanism is for, or what its function is, by identifying to an *adaptive problem* it solves. The algorithmic level would describe the form of algorithm that transforms inputs to outputs. And the implementational level would describe how such computations are physically instantiated in our brains (or perhaps another physical substance).

What the #@%! Is an Adaptive Problem?

Adaptive problems refer to the challenges and obstacles to survival and reproduction that our ancestors faced in their ancestral environments. Adaptive problems are challenges that would have needed to be reliably overcome generation after generation throughout our evolutionary history—often going back even before humans evolved. These problems encompass a range of fundamental concerns such as finding food, shelter, forming alliances, selecting mates, caring for offspring, avoiding predators, and navigating social hierarchies.

Evolutionary Psychology posits that solving these adaptive problems would require a sophisticated array of psychological mechanisms, each of which are designed to solve a particular subproblem. EP attempts to reverse engineer these psychological programs by conducting task analyses—a step by step set of instructions for all the little subgoals that need to be met in order to solve the broader adaptive problem. These task analyses can lead to novel predictions about the features of the human mind that can be tested in empirical studies. Ultimately, using adaptive problems to guide our study of human psychology sheds light on the evolved traits and behaviors that continue to shape human cognition and actions in contemporary societies.

7 Genetics and Personality

The winnowing effects of evolutionary processes tend to remove genetic variation from populations over generations. Because of this, any personality differences that are rooted in genetic differences would require us to develop an understanding of how or why evolutionary processes maintained this variation. But if personality isn't heritable then there is no real puzzle to solve.

So, is personality heritable? In this chapter, we'll explore behavioral genetics research that provides some insight into this question. In short, we will see that all personality traits—and likely all psychological traits—have a nontrivial genetic component. I'll present evidence from twin studies and molecular genetics that support this point. Finally, we'll consider how this genetic variation might be maintained by evolutionary processes.

Notes of Caution

Some people view talking about genetic influences to be controversial or taboo. It shouldn't be. If we want to understand how and why people are different or similar from one another, we must examine how genetics contributes to those differences. Part of the reluctance to examine the role of genetics might come down to some basic misunderstandings about what genetic differences mean and how genes influence traits. Let's try to address a couple of these issues before we dive into research on the heritability of personality.

Genetic Differences in Context

This chapter necessarily puts a lot of focus on genetic differences. But it is important to put these genetic differences in context. According to the [National Human Genome Research Institute \(2025\)](#), humans are on average ~99.6% genetically identical! We are overwhelmingly the same due to our very recent (in evolutionary terms) common ancestry. Less than 0.4% of the human genome (i.e., all of the DNA that codes for a human) differs between people.

The tiny fraction of genes in the genome that varies from person to person are referred to as *polymorphic genes*. These genes are what is being studied in the behavioral genetics research we are discussing in this chapter. Although it is just a tiny fraction of the genome, we'll see that this variation may have significant influences on human physiology, behavior, and cognition.

In behavior genetics, the outcomes being studied are referred to as phenotypes. It is important to note that because we are studying the small (<1%) part of the genome that *varies* across individuals that behavior genetics can only inform our understanding of *variation* in phenotypes; it does not tell us why an average level of the phenotype is observed. For example, behavior genetics does not help explain why the average height of a male in the USA is 5 foot 9 inches, but it can provide insight into the genetic and environmental contributions to differences in height between individuals.

Genes Are Probabilistic, Not Deterministic!

Although we will see that genes contribute a great deal to population variation in personality traits (and all other psychological traits that we have studied so far), we must appreciate that genes are not deterministic. Our genetics may predispose us towards certain characteristics, but only in a probabilistic sense via a complicated and cascading causal pathway. That means that genes only increase the odds of certain traits or outcomes—they don't cause them directly with any certainty in most cases. Further, the effects of genes on phenotypes depends on the environment: genes can have different effects under different environmental conditions. By recognizing that genes have probabilistic rather than deterministic effects, we can avoid falling into thinking that outcomes are inevitable and instead focus on how we might change environments to affect the likelihood of different outcomes.

Nature and Nurture(s)

Humans have debated for centuries whether behavioral differences are the result of nature or nurture. As with most debates, the extremes of either side are almost certainly incorrect. Behavioral differences are not likely to be completely due to differences in genetics; nor are they likely to be totally because of differences in upbringing. It is more fruitful to think of personality differences as arising from an interaction of genetic factors and environmental factors.

Moreover, the environment is not just one thing. We can distinguish between shared environmental features and non-shared environmental features.

When people think of the nurture side of the debate, they are probably thinking of the *shared environment*. The shared environment is made up of features of the environment in a person's upbringing that is shared across other family members (e.g., siblings). Common examples are things like parenting styles, socioeconomic status, or the number of books on the shelves in the house where a person grows up.

The *non-shared environment*, in contrast, is made up of all the experiences across development that are unique or different across family members. Some examples of the non-shared environment would be friend groups, different pop culture influences, random life events (e.g.,

accidents, illnesses, injuries). But even things like parenting styles or socioeconomic status or books on the shelf could become part of the nonshared environment if different siblings are exposed to different levels or types of these factors. For example, if parents read a parenting book that changed their parenting style between their first born and second born child, then parenting style could be part of the nonshared environment for each child. Finally, the non-shared environment also includes completely random noise and mutations that arise through development, as well as any errors in our ability to measure the trait of interest (i.e., differences in reliability and validity). In psychological research, where most of our measurements can be expected to contain a great deal of error, this is a big deal.

Until about 100 years ago, the nature-nurture debate was dependent on anecdotal evidence. There was not systematic research to inform the debate about the roles of genes and environments. But twin studies changed that.

Estimating the Contributions of Genetics and Environments with Twin Studies

Over the past 100 years or so, researchers have taken advantage of the nature of twins to estimate the influence that genetics and environments—both shared and non-shared—have on a variety of physiological and psychological traits. Twins are essentially a natural experiment that allows us to examine the roles of nature and nurture. What is it about twins that allows this?

Two Types of Twins

You are probably aware that there are two types of twins: identical and fraternal. These lay terms for twin types correspond to twins that are *monozygotic* and *dizygotic*, respectively. Unless twins are separated at birth or sometime in their childhood, we can assume that they experience pretty much the same “nurture” growing up. That is, they have the same parents with the same parenting styles, socioeconomic status, and number of books on the shelf. This is true of both monozygotic twins and dizygotic twins. However, the two types of twins both have very different genetic similarities because of the way the twins are conceived. Most of the research on twins is conducted on twins raised in the same home¹.

Monozygotic twins are the result of one ovum that has been fertilized by one sperm, which then happens to split into two zygotes. This creates two (mostly) genetically identical embryos. Roughly two-thirds of monozygotic twins also share a placental sac (i.e., chorion) in the womb. That means they are genetically identical, *and* most monozygotic twins are exposed to the same environmental factors throughout their embryonic development. So, they only thing

¹Studies based on twins raised in the same home are often referred to as “Classical Twin Studies”.

that can contribute to their differences are post-birth non-shared environmental differences (e.g., randomness, different friends, different classrooms).

Dizygotic twins, on the other hand, result from two different ovum being fertilized by two different sperm, creating two different zygotes. They only share 50% of the DNA at polymorphic gene loci: of the 0.1% of the genome that varies between people, they have the same genes for half of those differences and different genes for the other half. Which 50% of the polymorphic loci are shared or variable is itself random. That means they are no more genetically similar than the average pair of siblings who are not twins. The resulting embryos then develop in a separate chorion. So, their differences can arise from non-shared environments (e.g., hormone exposure in the womb) and genetic differences, but not from aspects of the shared environment (e.g., parenting) which is assumed to be the same for each twin.

These differences essentially make dizygotic twins analogous to a “control group” in a typical experiment. By comparing the similarity of monozygotic twins to the control group of dizygotic twins we can ask what the effect is of having the same genes (i.e., how heritable the trait is) on various phenotypes (e.g., height, religiosity, personality). Heritability is estimated by using Falconer’s formula.

Heritability and Falconer’s Formula

Heritability in twin studies refers to a very specific estimate: h^2 . This value is an estimate of the average amount of variance in a trait within the population being studied that can be attributed to variance in shared genetics. The heritability estimate we get from twin studies can lead to confusion if not interpreted correctly.

Because it is an estimate of the average contribution of genetics, it cannot be applied to individuals. So, if we learn that a trait is 50% heritable, it does not mean that your standing on that trait is 50% caused by your genes—it could be 30% or just about any other number. It only says that, on average, the contribution of genes to the trait is 50% across *all* individuals in the population.

Further, because the estimate only applies to the population sampled, the heritability of traits can differ across cultures and across time. The heritability of traits can be affected by many things. A classic example is that the heritability of height shrinks in environments where differences in malnutrition is common. In such environments, a person can have genes that predispose them to be tall, but malnutrition can stunt their growth. In places where there is plentiful food resources and low malnutrition in development, however, genetics contribute more to the population variation in height because nutritional differences between people can be smaller. Because heritability is specific to the population being studied, it is important to recognize that twin studies cannot be used to make inferences about differences between groups. Rather, they only inform our understanding of average contributions within a particular group and, as noted above, these average contributions still should not be applied to the individual.

In twin studies, heritability is estimated with *Falconer's Formula*: $2(r_{MZ} - r_{DZ}) = h^2$. In this formula, r_{MZ} is the correlation between monozygotic twin pairs on the trait of interest, and r_{DZ} is the correlation between dizygotic twin pairs on the same trait. Because the only difference between monozygotic twins and dizygotic twins, on average, is assumed to be genetic similarity, then twice the difference in their similarities is the estimated effect of genetic variation.

If we wanted to find the heritability of height, we would collect data on the heights of a large sample of monozygotic twins and a large sample of dizygotic twins. We would then compute the correlations between height for twin 1 and twin 2 the monozygotic twins and do the same for the dizygotic twins. Then we simply plug those values into the formula. For example, if researchers estimate the correlation between the heights of male monozygotic twins in Australia to be .84 and the correlation between the heights of male Australian dizygotic twins to be .45, the heritability estimate for height would be $h^2 = .78 = 2(.84 - .45)$. Thus, about 78% of the population variance in height can be attributed to variance in shared genetics.

But we must keep in mind that this doesn't mean that any given individual's height is 78% percent due to their genetics; nor does it mean that height in New Zealand or China is 78% heritable; nor even that the height of female Australians is 78% heritable. We only apply the heritability estimate to the population it was estimated in. We'd have to measure the heights of Australian female twins or Chinese twins to get their respective heritability estimates.

Heritability of Personality

Now that we understand the rationale behind estimating and interpreting heritability, we can discuss the heritability of personality traits. Across many studies with samples representing many populations, the heritability estimates for personality traits ranges from around 30-60%. The table below shows some representative heritability estimates from studies of the Big Five traits and Honesty-Humility.

Table 7.1: Representative examples heritability h^2 estimates for Big Five traits and Honesty-Humility

Trait	h^2
Extraversion	.55
Agreeableness	.41
Conscientiousness	.44
Neuroticism	.41
Openness	.34
<i>Honesty-Humility</i>	.40

Note: Big Five data from Jang et al. 1996; Honesty-Humility data from Lewis & Bates 2014. Don't take these as the absolute truth, just examples of what these estimates typically look

like.

Environmental Influences on Personality

If personality is about 30-60% heritable, what roles do the shared and non-shared environments play? To get the total environmental contribution we just subtract the estimated heritability from 100. So, if we use the estimate for the heritability of Extraversion in the table above, the environment explains about 45% of the variability in Extraversion on average (i.e., $1 - .55 = .45$). To parse the total environmental influence into shared environmental effects and non-shared environmental effects, we go back to our assumptions about similarities between monozygotic twins.

Because we assume that monozygotic twins are essentially identical in genes and the features of their shared environment (e.g., parenting, socioeconomic status, number of books in the house), then the only thing that can make them different is the non-shared environment. So, we just look at the difference between the correlation between dizygotic twins and a perfect correlation: $1 - r_{MZ}$. In the case of personality, the effect of the non-shared environment is estimated to be between 30-50%.

It is important to note again that the non-shared environment can include things like random noise and even measurement error—which is especially an issue in measuring psychological traits like personality. So, things like peer groups, teachers, and life experiences don't necessarily contribute the full 30-50% on average. A lot of the variation attributed to the non-shared environment could just be noise and measurement error.

If the non-shared environment explains about 30-50% of personality on average, and shared genetics explain 30-60% on average, that leaves only 5-10% left to explain on average. That remainder can be attributed to the shared environments. It can also be calculated by taking the difference between the monozygotic twin correlation and the heritability estimate: $r_{mz} - h^2$. This is because we assume that the only things that could make twins similar and explain the correlation between twins are genetics and shared environments. By subtracting out the estimated genetic effects, we are left with an estimate of the contribution by the shared environment.

In sum, it would appear from the research to date that there is a substantial genetic contribution to personality differences on average. And while the combined contribution of environments to personality variation is quite large as well, only the non-shared environment seems to contribute very much while the shared environment contributes relatively little. This is in striking contrast to what most peoples' intuitions are about the roles of nature and nurture in personality.

The Genetic Influence on Personality Changes Through Development

Twin studies of personality traits have been going on for decades now, so we can estimate heritability at different ages. In meta-analysis of twin studies across different ages, Briley and Tucker-Drob (2017) showed that the influence of genetics on personality declines throughout the lifespan, while the influence of the non-shared environment increases over time. They estimated the contribution of the shared environment to be essentially zero throughout the lifespan. Genetic influences account for most of the variation in personality from birth to about 25 years of age, when the non-shared environment contribution is equal. After about age 25, the non-shared environment overtakes genetic contributions to personality variation.

This is striking because some psychological traits, such as intelligence, show a different pattern whereby the effect of genetics increase throughout development. To me, this suggests that personality unfolds over time through differences in life experiences as we learn to adapt to the unique circumstances that we find ourselves in.

Again, these results only apply to a population of people, not to individuals within the population. For some people, the genetic influences may start off weaker and overtake the non-shared environment; for others the genetic influences may remain stronger than the non-shared environment over time; and others may exhibit different patterns of influence over development. Perhaps one day it will be possible to understand these individual differences in the trajectory of genetic and environmental influences on personality development.

Getting at the Genes Through Molecular Genetics

While twin studies allow us to estimate the extent to which variability in traits is related to variability in genetics, they don't measure any genes directly. So twin studies don't give us indication of how the genetics of personality work, or how many genes are associated with a given personality traits. To learn about the genetics of personality, we need to measure genes directly.

That's where molecular genetics comes in. There are two major forms of molecular genetics in personality research. The first is candidate gene research, which is hypothesis driven. The second is genome wide association studies, which is largely hypothesis free. We'll explore more about what this means each below.

Let's quickly review some key genetic concepts. *Genes* are segments of DNA that code for how to make a specific protein, and each gene is made up of multiple (often thousands) of *alleles*. There are also large regions of DNA that are referred to as either non-coding or intergenic regions, which do not directly code for proteins, but are made up of these same alleles. However, recent evidence from molecular genetic studies indicate that these non-coding regions may be among the most relevant for individual differences in traits like personality, likely because of their effect on the degree to which a gene is expressed (i.e., how much of a protein the gene

is making). Alleles may vary across individuals, and these variations can affect individual differences in traits or characteristics. Genetics research focuses on these *polymorphic* variants that vary across individuals.

The Shortcomings of Candidate Gene Research

Candidate gene research of personality was popular in the late 1990s, but they've been out of fashion for several years now. So why are we talking about them? While this sort of research basically failed to find reliable associations between specific genes and personality traits, I think their failures provide important context for understanding the complexity of the genetics of personality.

Candidate gene studies attempted to examine how specific genes relate to personality traits. Until recently, collecting DNA and sequencing large portions of the genome was time consuming and expensive. So, researchers relied on theory to guide their selection of which genes to look at. Specifically, researchers looked at small numbers of specific genes that previous research suggested may be related to various biological functions that could affect personality.

For example, researchers learned that a polymorphism on the *5-HTTLPR* was involved in serotonin transport. Because evidence suggested that serotonin was related to mood and feelings of happiness, researchers hypothesized that variation in serotonin-related genes would predict differences in personality traits like Neuroticism and Emotional Stability.

Many initial studies found the predicted association between the polymorphism and personality. However, over time the picture became much less clear because many studies did not find the link between the *5-HTTLPR* gene polymorphism and personality. This discrepancy may be because early studies were based on small samples that are likely to lead to false positives (i.e., finding an association between things that are actually unrelated). Studies with much larger samples have failed to find any reliable relationship between the polymorphism on the *5-HTTLPR* gene on Neuroticism or any other personality traits (e.g., Terracciano et al. 2009).

This pattern of early initial support followed by failures to replicate and inconclusive findings is very common in the candidate gene research on personality variation. Researchers now mostly recognize that the candidate gene approach is flawed because it assumes that personality trait variation will match closely with biological differences that can be explained by few genes with relatively large effects. But if many genes affect personality traits and each gene has only a small effect, then collecting data on only a small number of genes will never give us the whole picture.

The Promise of Genome Wide Association Studies (GWAS)

As researchers were realizing that the candidate gene approach was likely too narrow, genetic sequencing was also getting substantially more efficient and less expensive. Researchers could now measure larger portions of the genome in larger groups of people. This technological advancement fostered the emergence of genome wide association studies (GWAS; pronounced “jee-was”), which provide an alternative way to gain insight into the underlying genetics of personality.

The GWAS approach to prospecting for genes is hypothesis free—it doesn’t rely on researchers making guesses about which small number of genes to look at. Instead, GWAS studies measure larger portions of the ~.4% genome that differs between people (remember that ~99.6% of genome is identical from person to person) and explore which genes are associated with personality. If we use mining as an analogy, candidate gene studies are like using a pickaxe to start excavating in a specific place where someone told you gold might be, while GWAS research is like using explosives to blow a hole in the mountain to find the gold more quickly.

GWAS researchers measure portions of the genome that differ between individuals. This often involved studying the effects of *single-nucleotide polymorphisms*. Single nucleotide polymorphisms, which are often referred to as SNPs (often pronounced, “snips”) contain different alleles across different individuals within a population. If your genome is like a book, the alleles are the individual letters that make up different words. When these letters differ from person to person, we have a SNP. By examining variation in many SNPs, research can see how variability of SNPs on specific genes relates to variation in personality traits. Importantly, most GWAS researchers still don’t measure the entire genome, especially SNPs that are very rare, because it’s still not quite cost effective.

GWAS studies have been going on since at least 2008. Early on, sample sizes were still relatively small but in last decade most GWAS studies have been based on samples with hundreds of thousands or even millions of people. As the sample sizes have gotten bigger, researchers have been able to identify more and more genetic variants that are reliably associated with personality traits.

The main finding of GWAS research to date is that personality traits—and really any complex trait—are *extremely polygenic*. That means there are many, many genes that are associated with any given trait. Typically, dozens or hundreds or thousands of genes are reliably associated with any given trait, and each gene has only a very small effect size. According to a GWAS study of over 600,000 people, the estimated SNP-based heritability for Big Five traits is roughly 10–16% with contributions from over 1,000 genes, most of which are novel variants (Schwaba et al., 2025).

So, it doesn’t really make sense to talk about which specific genes may cause which personality traits. In fact, many genes are associated with more than one personality trait! Even though personality traits are roughly independent when we measure them using popular scales, the underlying genetics of each personality trait are correlated, meaning that many of the genes

that are associated with one personality trait are also associated with other personality traits (Lo et al. 2017).

Going forward, GWAS studies provide a promising avenue from which to identify genes that are reliably associated with traits. This will provide a more reliable foundation from which researchers can focus more directly on the biological functions of those sets of genes that predict personality. Ultimately, this may give us a firmer understanding of the underlying biology of personality. One thing appears certain: genetics won't map onto complex traits like personality in simple ways.

Missing Heritability(?)

Because twin studies estimate that population variation in personality is somewhere between 30 and 60% explained by genetics, we should eventually be able measure all the genetic variation in genomes within a population and explain 30-60% of the variation in personality using those genes. But so far, our best heritability estimates for personality variation in GWAS studies only explain about at most 16% of the population variation in personality traits. This discrepancy between what twin studies say *should* be explained by genes and what we can *actually* explain by measuring genes has been dubbed the “missing heritability” problem. Where might this missing heritability be?

One potential explanation has to do with the fact that GWAS studies are currently not based on the entire portion of the genome that varies—just a small portion of it where some *common* mutations are found. But there are other genes that vary from person to person because of *rare* mutations. It's likely that these rare mutations also have effects on complex traits like personality. Some recent work is beginning to find evidence for contributions from these rare genes (Wu et al., 2024), and future studies that rely on whole-genome sequencing ultimately fill in the missing heritability. Time will tell².

An alternative explanation is that twin studies may be overestimating heritability. As we discussed earlier, twin studies only provide accurate estimation under some very stringent assumptions that may not be completely satisfied by real world twins. For example, twin studies assume that monozygotic and (same sex) dizygotic twins have the same similarity of environments, but if monozygotic twins' environments are more similar, then this will lead to overestimation of heritability. Additionally, if genes and environments are correlated (a phenomenon we'll explore more below), then heritability will be overestimated. In short, it's possible that the heritability is not “missing” to the degree that current twin estimates would imply.

²I updated this Fall 2025. I should update this each semester or year to see what has changed. Bug me if I haven't!

Limitations of GWAS

Although GWAS research is delivering useful insights and promises to continue to do so, there are still some major limitations that need to be addressed. It is important to recognize these limitations so you can think critically about GWAS findings. Doing so may be necessary to make informed personal and policy decisions in the not-too-distant future with the developments happening in personalized genetic testing for medical interventions and gene-editing technology (e.g., CRISPR). Let's examine the most pressing limitations: population stratification, lack of diversity, ascertainment bias, gene-environment correlations.

Population Stratification

Population stratification refers to the presence of systematic differences in genetic ancestry or allele frequencies among subpopulations within a larger population. In GWAS, this can be a confounding factor because it can lead to false associations between genetic variants and traits if not properly accounted for.

Consider a genetic variant that shows population stratification to the extent that it has become more common due to recent genetic ancestry in western Europe and is less common in the United States. Across these geographic regions, there are going to be both these kinds of genetic differences due to population stratification in addition to culturally driven differences in certain traits. For example, soccer is the dominant sport in western Europe but not the United States. If we then conducted a GWAS of soccer playing ability using participants from both regions, and did not control for population stratification, we would falsely conclude that the genetic variant that varies due to ancestry causes increased soccer ability. This is a false conclusion because these genetic variants are merely *confounded* with soccer playing ability, which would be primarily influenced by cultural differences across these two groups.

To ensure that the observed genetic associations are indeed related to the trait being studied, rather than serving as a confound for differences in genetic ancestry, researchers typically focus on just one population that has recent common ancestry and attempt to control for genetic differences due to population stratification in their statistical analyses. But this approach can still leave subtle differences in population stratification that bias estimates of genetic effects. A recent study suggests that population stratification may not be a major issue in personality genomics, at least in populations with recent European ancestry (Schwaba et al., 2025).

Lack of Diversity

Researchers typically try to address confounding from population stratification by focusing on relatively homogenous samples with recent shared ancestry. In practice this means that most of the samples collected in GWAS research are composed predominantly of people with recent

European ancestry. This may be largely because such samples have historically been easier for researchers, who are predominantly working in the USA and UK, to access.

This lack of diversity can severely limit the generalizability of GWAS findings to other populations with different recent ancestry. Genes may have different associations with outcomes (e.g., disease, personality) across different populations, so the lack of generalizability may perpetuate inequalities. From the limited tests that we do have examining how well findings from one population generalize to another (e.g., Gupta et al., 2024; Schwaba et al., 2025), it seems it seems like there are important differences in the genetics of personality across human populations. Representation has slowly improved in recent years, but over 80% of GWAS samples are still conducted in populations with recent European ancestry, even though such populations represent only 16% of the world population (Mills & Rahal, 2019).

Ascertainment Bias

Ascertainment bias occurs when the selection of participants for a study is not random or representative of the larger population. In the context of GWAS, ascertainment bias can arise when the selection of individuals is influenced by certain characteristics or conditions related to the trait being studied. This bias can skew the results and make them less applicable to the broader population.

For example, if agreeing to give DNA to anonymous faceless researchers or companies for research is in part determined by personality (e.g., agreeableness, neuroticism, openness), then this can bias estimates of the associations between genes and personality. Further, even within the pool of people who are willing to give their genetic data to research, there is probably still non-random genetic influences on whether people are willing to answer long lists of questions in psychological scales. This can further bias the results of GWAS studies, potentially leading to finding spurious correlations between genes and traits. Recent examinations of ascertainment bias find that this may be a significant issue because genetics predicts both personality and survey participation (Schwaba et al., 2025).

Gene-Environment Correlations

Gene-environment correlations refer to the interplay between an individual's genetic makeup and their environment. These correlations can be passive, evocative, or active. In the context of GWAS, researchers recognize that an individual's genetic predispositions can influence the environments they are exposed to or the choices they make, which in turn can affect their personality. So, it becomes difficult to disentangle the effects of genetics and environments because environments are related to genetics. A recent large-scale GWAS study used sophisticated statistical methods to examine the extent of bias introduced by gene-environment

correlations but found very little evidence that this is a problem for genetic inference when it comes to personality traits³ (Schwaba et al., 2025).

Conclusions About Genetics and Personality Heritability

This chapter examined whether and how genetics are related to personality. We learned that the differences in genetic similarity between monozygotic and dizygotic twins provide a natural experiment for researchers to estimate the relative contribution of genes, shared environments, and nonshared environments to personality variation. These so-called twin studies suggest that somewhere between 30 and 60% of the variation in personality traits within a population can be attributed to similarity in genetics. The remaining 40-70% of the variation in personality can be attributed to environmental influences. But very little of that variation can be attributed to shared environments—most of the variation in personality is attributed to the non-shared environment. Genome wide association studies allow us to explore the genetic underpinnings of heritability. They show that personality traits are extremely polygenic, meaning that thousands of genes influence personality in very small ways. Further, GWAS studies suggest that many supposedly distinct personality traits have overlapping genetic underpinnings. Although there are still many issues to work out before we can hope to understand the complex genetic underpinnings of personality, there is lots of promise. Ultimately, it is important to recognize that even though humans are 99.9% genetically identical, our small genetic differences may still play an important role in driving the diversification of human personality.

Potential Evolutionary Explanations

Let's return now to the evolutionary puzzle of personality we started this chapter with: if personality is heritable, why (or how) has it been maintained by evolutionary processes that tend to remove genetic and phenotypic variability from populations over time? Personality does indeed seem to be heritable, so an explanation for why is needed. There are several potential explanations, which we will explore below. In general, there are two types of explanations: functional and non-functional.

Functional Explanations

Fluctuating/balancing Selection

If some level of trait is always better for survival and reproductive success, then selection will tend to remove variation around that optimal value. However, if the fitness benefits of

³But note that similar investigations of other traits, like cognitive ability and educational attainment, found evidence of bias from gene-environment correlations in estimates of SNP-based heritability of these traits.

a trait differ across environments (e.g., cultural contexts) over time, then there is no optimal level for evolution to zero in on. This is called *fluctuating* or *balancing* selection because the fitness of certain traits fluctuates. For example, it could be that Extraversion is a trait that is beneficial in environments where there are low risks of getting sick from interacting with people, extraversion may be harmful on average to survival and reproduction in environments with increased risk of contracting illnesses from interactions (e.g., pathogen rich environments). Daniel Nettle, an evolutionary psychologist, used cost-benefit analysis to examine the potential fitness impacts of different personality traits across variable environments and concluded that fluctuating selection could maintain the personality variation that we observe today (Nettle, 2006). This is one of the most popular hypotheses for the evolution of personality traits, but as we will see later, recent genomic evidence casts doubt on this story.

Frequency Dependent Selection

Another way personality variation could be maintained is through a process called frequency dependent selection. Under frequency dependent selection, the evolutionary fitness of a trait depends on the relative proportions of the trait in the population. Some traits or behavioral strategies are only good for fitness when others in the population have different traits or are playing different behavioral strategies. Some psychologists have suggested that traits like sociopathy—which is characterized by callouses, lack of empathy, and willingness to cheat others—might be maintained via frequency dependent selection. The argument is that people with sociopathic traits can only effectively manipulate or exploit other people who do not have those traits because other sociopaths are more difficult to manipulate. That means that sociopathic traits would only be maintained at low frequencies in the population. I don't know of any clear tests of this hypothesis, but it is possible.

Non-Functional Explanations

Selective Neutrality

Selective neutrality means that a trait is neutral with respect to its susceptibility to evolutionary selection pressures. Essentially, a trait would have to be completely unassociated with survival and reproductive success in order to be selectively neutral. This explanation seems particularly unlikely for personality traits given that there are many obvious (and probably many nonobvious) ways that personality traits can influence survival and mating success. Humans even select mates based on their personality traits, so it would be very unlikely that personality traits are invisible to selection.

Mutation-Selection Balance

Another non-functional alternative is that personality traits are not adaptations in themselves but reflect variation in adaptations that arise through random mutations. Because each generation has its own unique set of random, mostly deleterious mutations that arise, selection processes can never completely remove genetic variation from the population. (and in fact, these random mutations are the raw ingredients of evolution). This means that personality variation might just be a product of noise from mutations that make subtle differences in behavior and cognition between people. But these differences are important to other humans for predicting how people will act and deciding who to interact with, so it is possible that the difference detecting mechanisms to think about personality and develop lexical concepts to describe differences may itself be an adaptation (cultural and/or genic).

As of updating this chapter in Fall 2025, there have been a few advances in genomics that can help us begin to rule out some these potential hypotheses about the evolution and maintenance of personality variation in humans. New large-scale GWAS research (e.g., Schwaba et al., 2025) suggests that the DNA influences on Big Five traits mostly sit in brain-neuron genes that don't tolerate damaging changes. That pattern is a hallmark of purifying (negative) selection—harmful changes tend to get weeded out. This argues against balancing/fluctuating selection as the main genome-wide reason personality differences persist and fits better with mutation–selection balance: new small mutations appear each generation, the worst are removed, and lots of tiny effects remain. In a recent review of the genomic evidence to date, the author concluded, “there is no evidence that balancing selection has substantively shaped complex traits, and strong evidence that it has not” (Zietsch, 2024, p.1).

Recent genomic studies sharpen what's plausible and what's not, but they don't crown a single winner. They do, however, shift the default toward “non-functional” baselines (mutation–selection balance) with room for local functional effects in specific contexts. Ultimately, more research is needed to understand the genetics of personality and to explain why heritable differences in personality have persisted throughout human evolution. But this puzzle highlights the importance of examining any phenomenon from both ultimate and proximate research programs. The ultimate research that explores the evolution of personality is complimented by the proximate research that explores the contemporary genetics of personality.

References

- Briley, D. A., & Tucker-Drob, E. M. (2017). Comparing the developmental genetics of cognition and personality over the life span. *Journal of Personality*, 85(1), 51-64.
- Gupta, P., Galimberti, M., Liu, Y., Beck, S., Wingo, A., Wingo, T., ... & Levey, D. F. (2024). A genome-wide investigation into the underlying genetic architecture of personality traits and overlap with psychopathology. *Nature human behaviour*, 8(11), 2235-2249.

- Lewis, G. J., & Bates, T. C. (2014). How genes influence personality: Evidence from multifacet twin analyses of the HEXACO dimensions. *Journal of Research in Personality*, 51, 9-17.
- Lo, M. T., Hinds, D. A., Tung, J. Y., Franz, C., Fan, C. C., Wang, Y., ... & Chen, C. H. (2017). Genome-wide analyses for personality traits identify six genomic loci and show correlations with psychiatric disorders. *Nature genetics*, 49(1), 152-156.
- Mills, M. C., & Rahal, C. (2019). A scientometric review of genome-wide association studies. *Communications biology*, 2(1), 9.
- Nettle, D. (2006). The evolution of personality variation in humans and other animals. *American Psychologist*, 61(6), 622.
- Terracciano, A., Balaci, L., Thayer, J., Scally, M., Kokinos, S., Ferrucci, L., ... & Costa Jr, P. T. (2009). Variants of the serotonin transporter gene and NEO-PI-R Neuroticism: No association in the BLSA and SardiNIA samples. *American Journal of Medical Genetics Part B: Neuropsychiatric Genetics*, 150(8), 1070-107
- Jang, K. L., Livesley, W. J., & Vernon, P. A. (1996). Heritability of the big five personality dimensions and their facets: A twin study. *Journal of personality*, 64(3), 577-592.
- Wu, X. R., Li, Z. Y., Yang, L., Liu, Y., Fei, C. J., Deng, Y. T., ... & Yu, J. T. (2025). Large-scale exome sequencing identified 18 novel genes for neuroticism in 394,005 UK-based individuals. *Nature Human Behaviour*, 9(2), 406-419.
- Zietsch, B. P. (2024). Genomic findings and their implications for the evolutionary social sciences. *Evolution and Human Behavior*, 45(4), 106596.

Acknowledgements:

Huge thanks to Dr. Andrew Grotzinger, assistant professor at The University of Colorado Boulder, for providing extensive feedback on this chapter. Also, thanks to Dr. Ted Schwaba for providing feedback on interpretations of his 2025 study.

8 Physiology and Personality

In the last chapter, we learned that genetics explains a significant portion of variation in personality and that the genetics of personality are far from simple. There are many, many genes that influence personality—and any single gene has only a very tiny effect.

The effects of genetics on personality must be mediated by our physiology. Genes are simply rough instructions for making proteins that build our brain and body¹. The resulting subtle differences in our physiology contribute to personality differences. Put another way, physiological differences in brains and bodies are what links genetic differences to personality differences.

In this chapter, we'll begin to explore how physiology relates to personality. Unfortunately, it will soon become clear that there are usually not easy or simple answers when it comes to how differences in brains and bodies are related to average differences in thoughts, feelings, and behavior. But I think it's important to recognize this complexity to avoid falling for overly simplistic explanations that boil personality differences down to simple physiological factors. Personality traits are complex phenomenon, and we should probably not expect that their physiological underpinnings are simple or straightforward.

We'll start with brains before moving to our bodies more generally. But let's not mistake this topic separation for biological reality. Minds, brains, and bodies are not distinct things. Our “mind” is a construct we use to describe the set of operations performed by our brain. Our mind is a product of our brain, and our brain is part of our body. Ultimately, explaining how physiology relates to personality will require understanding our bodies as a whole unit with interacting parts.

Brains

What Is a Brain?

Our brains are a highly sophisticated web of interconnected *neurons*. Neurons are a specialized form of nerve cell. They are the primary functional unit in the brain. The average human brain has 86 billion neurons (Herculano-Houzel, 2012). Each neuron can be connected to around 10,000 other neurons. That means that there are potentially trillions of neural connections

¹Remember our mantra: genes are probabilistic, not deterministic!

in the human brain! This complexity, and the resulting potential for an immense number of possible configurations of brain activity or “brain states”.

Neurons are responsible for receiving, transmitting, storing, and performing computations on sensory information from the external world, as well as information produced in the brain. They do this by communicating via their connections to other neurons. Neurons communicate with other neurons primarily using electrochemical signals called action potentials and chemical messengers called neurotransmitters. All our thoughts, feelings, motivations, behaviors, and memories are products of this neuronal communication.

Neurons are generally composed of a cell body with dendrites, an axon lined with myelin, and an axon terminal. Dendrites receive signals from other cells and transmit information as an electrical signal down the axon to the axon terminal. This transmission of energy from the dendrites through to the axon terminal is often referred to as the neuron “firing”. The axon terminal contains *vesicles*, which are capsules that store neurotransmitters (imagine a gelatin pill capsule that contains powdered medicine) until it is ready to release them.

When a neuron fires, an electrochemical signal travels to the axon terminal and causes vesicles to release neurotransmitters into a space between neurons called the *synapse*. Once in the synapse, neurotransmitters then bind to *receptors* on other neurons to cause them to fire their own action potential. A signal-sending neuron is referred to as a *presynaptic* neuron, and the receiving neuron is referred to as *postsynaptic* neurons. Neurotransmitters and receptors exhibit a “lock-and-key” style relationship, whereby the specific molecular structure of neurotransmitters is a perfect fit for specific receptors.

This electrochemical communication between neurons and larger brain networks is the primary way that computations happen the brain. Because of this, differences in neuronal communication may help us understand how average individual differences in thoughts, feelings, and behavior (i.e., personality) arise.

Neurotransmitters and Personality

There are over 100 types of neurotransmitters. But research in personality has explored on only a small number of these. In this chapter we will focus on two neurotransmitters that have been the subjects of most personality research: serotonin and dopamine.

It’s tempting to present individual types of neurotransmitters as having specific and narrow roles in neuronal communication. But reality is much more complex. Neurotransmitters actually have many diverse functions and are typically widely distributed throughout the brain and even nervous system.

Serotonin, for example, is involved in a wide variety of phenomena including sleep, mood, digestion and nausea, wound healing, and sexual arousal. Serotonin is linked to depression and anxiety (Stein & Stahl, 2000). Drugs like MDMA create feelings of euphoria by triggering your neurons in particular circuits in your brain to release large amounts of serotonin.

Dopamine plays an important role in movement, pleasurable feelings, motivation and addiction, and emotional arousal. Too little dopamine is linked with Parkinson's disease, which is characterized by uncontrollable body movements and shaking. Too much dopamine, in contrast, is associated with schizophrenia which is characterized by hallucinations, delusions, slow movement, and lack of motivation.

Considering the functions of these neurotransmitters, we could make predictions about their relations to personality. For example, we could predict that serotonin would be linked to the personality trait Neuroticism. And perhaps dopamine could be expected to be related to Openness or Extraversion. However, there is little evidence for reliable relationships between these neurotransmitters and personality.

Limited Evidence for Links Between Neurotransmitters and Personality

The strongest evidence for links between neurotransmitters and personality comes from clinical trials of pharmaceutical treatments for depression and anxiety. The first-line treatments for depression and anxiety include a class of drugs called selective serotonin reuptake inhibitors (SSRIs). As the name implies, SSRIs work by preventing presynaptic neurons from taking back serotonin that is floating in the synapse, which allows the postsynaptic neuron's receptors to have more time to bind with the serotonin.

Some research suggests that patients who are assigned to take SSRIs for treatment of depression or anxiety tend to score lower on self-report measures of Neuroticism in post-treatment questionnaires than in the pre-treatment questionnaires (Noordhof et al., 2019). This provides some tentative evidence that differences in serotonin are linked to differences Neuroticism. However, it isn't clear whether differences in serotonin can account for differences in Neuroticism (or other personality traits) in people who are not seeking treatment for depression or anxiety.

Complicated Relationships May Obscure Associations

The lack of evidence for clear links between neurotransmitters and personality variation may simply be because there are very few simple or straightforward links between neurotransmitters and personality. The effects of neurotransmitters depend on a multitude of factors. Simply having more or less of a given neurotransmitter does not necessarily mean that it will have more effects. Figure 8.1 highlights some of the many aspects of neuronal communication that could lead to individual differences in the computations that produce behavior, thoughts, and feelings.




Figure 8.1: Potential sources of individual differences in neuronal communication and computation.

For example, there are differences in the number of postsynaptic receptors available to take up the neurotransmitters. There are also differences in how long the neurotransmitters stay available in the synapse. This can be caused by differences in the speed with which presynaptic neurons reuptake free-floating neurotransmitters, as well as differences in the speed with which the enzymes that break down free-floating neurotransmitters. Even factors like individual differences in the space between neurons, vesicle size, or differences in the viscosity of the fluid that neurotransmitters are transported through could potentially lead to differences in the chemical computations underlying thoughts, feelings, and behavior.

Therefore, it seems unlikely to me that the complexity of individual differences in neurotransmitter communication will be captured by examining simple relationships between neurotransmitter levels and personality traits. Without measuring more of the relevant factors and accounting for the complex interactions among all these moving parts, it will be difficult to understand how neurotransmitters relate to personality. In order for differences in neurotransmitters to explain personality differences, changes to neurotransmitter levels (1) would need to be sufficient to produce changes in activation in large brain regions/circuits/networks (otherwise the effects would be micro and insignificant); and (2) would need to be sustained over a long enough period of time to exert a detectable effect on personality features (otherwise the effects would be transient). Ultimately, neurotransmitters are unlikely to be linked to personality in a clear 1-1 fashion. We will need much more sophisticated methods and larger sample sizes to study these complex relationships.

Brain Region Size, Connectedness, and Activity

Neuroscientists also look at higher levels of organization within the brain in search of the physiological underpinnings of personality. Some of the most common targets of research are differences in the *size*, *activity*, and *connectedness* of different brain regions to try to explain personality differences.

Underlying most of the research seems to be the assumption that if an area of the brain is bigger or more active or more connected, it will lead to a person's personality being more reflective of the types of thoughts, feelings, and behaviors that are associated with that brain area. For example, if the frontal lobe is associated with impulse control, planning, and decision making (i.e., "executive functioning") then people who have bigger, or more active, or more interconnected frontal lobes are expected to be more skilled at executive functioning, which may be predicted to relate to personality traits like conscientiousness.

This thinking goes back even before modern neuroscience. There have been many historical cases of brain damage changing a person's personality, which seem to support these assumptions. For example, Phineas Gage accidentally shot a railroad spike through the front of his brain, and his personality is reported to have changed from an easygoing, kind individual to a hot-tempered, impulsive individual. However, all historical accounts are subject to embellishment over time. To really test whether differences in the brain are predictive of differences in personality, we need to examine contemporary neuroscience research with appropriate experimental controls and comparison groups.

Modern neuroscience aims to systematically map the human brain and its functions by examining its structure and activity patterns. As a side goal, neuroscientists can gain insight into individual differences in brain structure and activity, which may inform personality research. Let's examine some of the main neuroscience methods and see what they may reveal about personality.

Brain Structure

Neuroscientists have developed variety of neuroimaging tools for examining brain structure. Indices of brain structure include the size of brain regions, the shape of brain regions, and the interconnectedness of brain regions. It is helpful to understand neuroimaging methodology to envision the research and aid in interpreting the results, so we will review some common neuroimaging methods here.

Differences in the size of brain areas can be measured with **Computerized Tomography (CT)** or computerized adaptive tomography (CAT) scans. CT scans are based on passing many x-rays through the brain and measuring how their signal weakens in order to visualize the brain. This provides low resolution images of the brain that can be used to assess the size and shape of brain regions.

Similarly, differences in the size and shape of brain regions can be measured with structural **Magnetic Resonance Imaging (MRI)**. MRI scans use a very strong magnet and send radiofrequency pulses through the brain and measure how protons throughout brain take to recover from these pulses. The denser the nervous tissue in a brain area, the larger this recovery signal. These methods provide higher resolution images that can be used to assess the size of brain regions.

Diffusion Tensor Imaging (DTI) provides further insights into the structural connectivity of the brain. It focuses on the movement of water molecules as they diffuse across groups of neurons that form white matter tracts. In the brain, water molecules tend to move more easily along the long the axons of neurons. DTI capitalizes on this tendency by measuring the direction and magnitude of water diffusion in multiple directions. This information allows researchers to map and visualize the brain's white matter tracts, which are bundles of neuron fibers that connect different brain regions.

What does neuroscience research using these methods to examine differences in the size, shape, and interconnectedness of brain regions reveal about personality? Unfortunately, the answer does not appear to be as clear as we might hope. There have been dozens of studies over the years, but each finds different inconsistent associations between indices of individual differences in brain structure and personality. Because neuroimaging studies are quite expensive to run, they typically rely on very small samples, so it is difficult to trust any one study.

A recent meta-analysis that systematically surveyed and aggregated the near totality of studies examining links between personality and brain structure reported that there is “nothing to see here”, meaning that there is no evidence for reliable links between gross brain structure differences and personality (Chen & Cali, 2022). Similarly, recent studies that examined over 1000 subjects each found mostly null or very small relationships between gross brain morphology and personality traits (Gray et al., 2018; Hyatt et al., 2021). In sum, there does not seem to be any strong evidence linking variation in gross brain structure, as measured by current methods, to variation in our current constructs of personality.

Brain Activity

Neuroscientists can also use neuroimaging tools to measure brain activity. These methods allow for spatial assessment of *where* activity is happening, quantitative assessment of *how much* activity is happening, and temporal assessment of *when* activity is happening. Let's review some details of these methods, so we can get a better understanding of how this research is done before evaluating what the research reveals about personality.

The most common method to examine brain activity is **Functional Magnetic Resonance Imaging (fMRI)** scans. fMRI looks at blood flow changes in the brain. When a brain area is activated, it consumes more oxygen, and blood flow increases to that area. fMRI detects these shifts in blood flow, helping us pinpoint the brain regions involved in specific tasks or

conditions. Moreover, fMRI can provide insight into how parts of the brain are interconnected by examining where blood flows in the brain.

Information about brain activity and connectivity can also be obtained via **Positron Emission Tomography (PET)**. PET uses a special kind of sugar that emits tiny particles (positrons) when it's used by active brain cells. When these particles collide, they create a signal that PET picks up, showing us which brain regions are most active at a given moment. A big advantage of PET is that researchers can examine the activity of specific chemical systems (e.g., the dopaminergic or serotonin systems).

Information about when and where *electrical* activity occurs in the brain can also be obtained efficiently via **Electroencephalography (EEG)** and **Magnetoencephalography (MEG)** methods. EEG uses small sensors placed on the scalp to detect the electrical signals generated by activity from large numbers of neurons across the skull. When brain cells communicate, they produce these electrical impulses that EEG can capture with fine temporal resolution, showing us the rapid changes in brain activity. MEG, on the other hand, is like a magnetic cousin to EEG. It detects the magnetic fields generated by neuronal activity. When neurons fire, they create tiny magnetic fields, and MEG sensors outside the head can detect and map these fields. Both EEG and MEG provide high-temporal resolution, allowing psychologists to observe brain activity with millisecond precision. By analyzing the patterns of these electrical and magnetic signals, researchers can explore the timing and coordination of brain activity during different mental processes, providing valuable insights into brain dynamics and how our thoughts and experiences unfold over time.

Together, these technologies help neuroscientists see which parts of the brain are active when we think, feel, or do something, giving us crucial information about how our brains function and how behaviors and thoughts are linked to brain activity. So, what does neuroscience research examining differences in brain activity reveal about personality? As was the case for brain structure, there have been many studies over the years, and each study tends to turn up somewhat different associations between indices of individual differences in brain function and personality. Again, many studies rely on relatively small samples, so it is difficult to trust any one study.

A meta-analysis by Kuper et al. (2019) aggregated 79 studies that looked at relationships between EEG activity and personality traits. In total, the studies represent over 5700 participants. The authors of the study reported that “less than 0.4% of the variance in extraversion and neuroticism could be explained by resting frontal asymmetry”, which is one type of EEG data. Similarly, a recent fMRI study that examined over 1000 subjects to test for associations between personality traits and brain activity in several regions during reward-based tasks found no reliable relationships (Hyatt et al., 2020). However, another study using machine learning methods to measure individual differences in the patterns of interconnected activation across many brain regions reliably predicted all Big Five traits, except extraversion (Cai et al., 2020). More recently, a meta-analysis of 17 studies identified reliable associations between neuroticism and resting-state brain activity in several brain regions (Lin et al., 2023). In sum,

the relatively inconclusive results of research linking brain size, activity, and connectedness to personality traits suggests that such links are complicated.

Tentative Conclusions About the State of Personality Neuroscience

At this point you may be feeling somewhat underwhelmed with the insights delivered by neuroscience methods to the study of personality. I am too. But many researchers are optimistic that research using larger samples and better methods may provide more insight (e.g., De Young et al., 2022).

A less optimistic possibility, however, is that differences in brain structure at the anatomical level are just not useful in explaining personality differences. Individual differences in any of these areas might have very different effects in different people. That is, the impacts of brain physiology on thoughts, feelings and behavior may be *heterogenous*. There may be no systematic or straightforward relations that are consistent across people. Additionally, the underlying assumption that the size, density, interconnectedness, or activity of brain regions can be expected to be related to the amount or strength of thoughts, feelings, behavior, or personality may need to be reexamined. Of course, this is ultimately an empirical question that will need to be resolved with more research rather than debate. I plan to update my beliefs on this as more data come in.

Another issue is that the personality traits like that Big Five and HEXACO, which are derived from lexical studies and factor analyses may not correspond to the underlying biology of our brains in any clear way. In order to build an understanding of whether and how our biology contributes to our personality, we may need to develop better theories of how personality traits work at the functional and algorithmic levels of analysis (c.f., Marr's levels of analysis from the Foundational Frameworks chapter). Only then can we know what we are looking for in at implementational level of analysis—that is, the neurons in our brain that conduct the computations that produce our mind.

Bodies

Now let's turn to aspects of our physiology outside of the brain. The aspects of our bodies that we will look at here interact with the nervous system in many ways and may have some influence on personality indirectly through their effects on our brain/mind. Let's start with the chemical messengers of the body: hormones.

The Role(s) of Hormones

Hormones, like neurotransmitters, are chemical messengers. But hormones are produced by endocrine glands such as the pituitary, thyroid, and adrenal glands, and then distributed

throughout the body. They travel through the bloodstream and target specific cells or organs, influencing a wide range of bodily functions and behaviors and do so for a relatively prolonged period of time. Hormones are crucial for coordinating bodily function and maintaining homeostasis, or balance, across the body's many organ systems.

The effects of hormones are most apparent during puberty, as hormones trigger the emergence of secondary sexual characteristics, such as breast development and body hair growth. They also influence psychological and emotional changes, affecting mood, behavior, and sexual drive.

Because hormones are also related to everyday behavior, thoughts, and feelings, personality researchers have reasoned that individual differences in hormone activity might help to explain personality differences. For example, differences in the amount or concentration of certain hormones may lead to average differences in behavior or thoughts or feelings. Also, because hormone levels fluctuate throughout the day, there could be differences between people in how much or how quickly their hormones change.

There are over 40 hormones in the human body. Any combination of hormones can be traveling through our bloodstream at any given time. I like to think of hormones more as some kind of thick soup with many ingredients—it is difficult to pick out the effects of any single ingredient (hormone) because it is their combination and balance that creates a flavor palette (behavior or moods).

Nonetheless, most research on the relationship between hormones and personality tends to focus on only one or two hormones at a time. This is probably because collecting hormones is costly and time consuming. Hormones are typically collected via blood, saliva, urine, and hair follicles. Most researchers do not have the ability to safely store or assay (analyze) hormones within their own labs, so they must pay for other companies to do so. It is getting less expensive over time, but the expensiveness of hormone research has historically led to relatively small samples in research studies.

We'll focus on the two most studied hormones when it comes to personality: testosterone and cortisol.

Cortisol is a steroid hormone, popularly known as the “stress hormone”. Cortisol can influence many bodily processes and organs because receptors for cortisol are found throughout the body. Cortisol is released in response to stress and helps regulate metabolism, manage stress levels, and support the body's fight-or-flight response.

Testosterone is an androgen that is typically present in higher levels in males than in females. During puberty, it drives the development of secondary sexual characteristics in males, such as the growth of facial and body hair, deepening of the voice, and muscle mass increase. This hormone is crucial for the regulation of energy metabolism, bone density, and the production of red blood cells. Additionally, individual differences in testosterone are associated with differences in aggression (Geniole et al., 2020).

Hormones and Personality

Given the roles of cortisol and testosterone, how might we expect them to be related to personality? You might predict that because cortisol is related to stress responses, and the personality trait of Neuroticism seems to capture differences in stress responses (e.g., depression, anxiety), then differences in cortisol levels will relate to differences in Neuroticism. Additionally, you might predict that testosterone will be related to Extraversion or Agreeableness because testosterone is linked to aggression and competition, and Extraversion or Agreeableness capture differences in related constructs of status-striving and domineering. This is essentially the predictions made by many researchers examining links between hormones and personality.

As seems to be the case with many lines of research on links between physiology and personality, there are dozens of studies in relatively small samples of participants—sometimes less than 20 people—with inconsistent or overhyped findings that are difficult to trust. Some studies find relationships between testosterone and Extraversion or cortisol and Neuroticism, while others do not.

The most convincing study that I've come across used data on testosterone and cortisol levels and Big Five personality traits for almost 4,000 people (Sundin et al., 2021). This study examined whether the two hormones predicted variation in any of the Big Five traits. They found that most of the associations between the two hormones and personality traits were not statistically significant. Only one correlation was statistically significant: testosterone was weakly negatively associated with Conscientiousness ($r = -.05$). Perhaps we can make sense of this by considering that testosterone differences are associated with differences in aggressive tendencies (Geniole et al., 2022). In many social contexts, aggressive behavior would be considered taboo and can have big negative consequences (e.g., fines, jailtime), so aggressive acts may lead someone to be perceived as low on conscientiousness (e.g., not caring about rules, not thinking of consequences). Maybe future research will provide more insight.

In general, it seems that there are not many clear, strong links between broad personality features and single hormones. Perhaps things will become clearer as research samples get larger. However, it is likely that the effects of hormones on personality depend on complex interactions among many hormones, so we will need very large samples to be able to reliably estimate these relationships.

Does Anything About Our Physiology Predict Personality?

At this point you may be a little frustrated. Isn't there anything about physiology that predicts personality? In this section, we will look at a theoretical framework that can give clearer predictions about features of our physiology may be predictive of personality variation. Specifically, we'll examine the *facultative calibration* hypothesis of personality variation.

Facultative Calibration!?

The phrase “facultative calibration” might appear a little daunting, so let’s break it down. The term “facultative” refers to traits or behaviors that are flexible or changeable. The term “calibration” is about fine-tuning or adjusting something to make it accurate or suitable for a particular purpose. In the case of personality, facultative calibration refers to adjusting tendencies in behaviors, thoughts, and feelings (i.e., personality) over development.

More specifically, the facultative calibration hypothesis posits that our personality development is guided (calibrated) by psychological mechanisms that calculate the costs and benefits of different behavioral strategies (e.g., aggressiveness, sociality, sexual promiscuity). The costs and benefits here are ultimately about evolutionary fitness—things that increase the chances of survival and reproduction are *benefits* from the perspective of our brains, and things that lessen the chances of survival and reproduction are *costs*. The costs and benefits of different strategies depend on many things, such as environmental factors, other peoples’ perceptions and preferences, cultural contexts—and our physiology. Essentially, the facultative calibration hypothesis assumes that our brain figures out what is a good personality trait for us to develop by assessing the likely costs and benefits of different traits, and investing in the development of traits that are less costly relative to their benefits.

At a basic level, the facultative calibration hypothesis leads the prediction that any physical trait that is related to the costs and benefits of a personality trait on average, will come to be associated with that trait. For example, if a physiological trait makes it easier to carry out a particular personality trait, then individuals possessing that physiological trait should be more likely to exhibit that trait.

Empirical Support for Facultative Calibration of Personality Traits

Thus far, empirical research on facultative calibration has focused on three physiological features: attractiveness, formidability (i.e., physical strength and fighting ability), and health. Each of these physiological characteristics can be linked to personality traits by considering whether they make particular personality traits more or less effective. Let’s explore some of these predictions and see whether they are supported by empirical data.

Attractiveness perceptions are ultimately an assessment of others’ mate value. If we were peahens, we would probably find a peacock with a big bushy and colorful tail to be very attractive because such a tail is a clue that the peacock has genes that are good at surviving and thriving. In humans, the things we find attractive also tie back clues about whether a person would be good to mix up our genes with to produce and offspring that could live to successfully mix up their genes with someone else in the particular ecological contexts that humans evolved in (Sugiyama, 2015). Because attractive people are desirable mates, they have more bargaining power in the mating world—and this often translates into better treatment in everyday life.

In language of costs and benefits, attractiveness can make it more likely that social relationships will be easier to make and more beneficial on average. Thus, we might expect that people who are more attractive tend to be more outgoing and maybe also more entitled. Indeed, there is evidence from a study with over 600 participants that people who are rated as more attractive by others tend to be a moderately higher in extraversion and lower in honesty-humility (Rodriguez & Lukaszewski, 2020).

Formidability is probably a crucial individual difference to our stone-age brains. Some of my own research suggests that our brains can accurately assess the strength of men in mere milliseconds (Durkee et al., 2018). This may be because differences in formidability—fighting ability and physical strength—would likely have contributed to inequality in access to resources, such as food, mates, and status throughout mammalian (and human) evolution. In aggressive contests over resources more formidable individuals have better odds of wining. Even though most of the social interactions we encounter in the social world don't depend on actual physical aggression, our brains may still largely interpret formidability as an important cue to someone's bargaining power.

In the language of costs and benefits, formidability lowers the costs of aggression and increases the probability of benefits from being aggressive or dominant in social bargaining situations. As such, we might expect that more formidable people are less agreeable, less fearful, and more entitled on average. Additionally, because formidability is an important component of (mainly men's) attractiveness, we might expect that formidability is also predictive of extraversion. Indeed, there is partial support for these predictions in large-sample studies: physically stronger people tend to score somewhat higher on measures of extraversion and fearfulness (e.g., Rodriguez et al., 2020; Manson et al., 2023).

Health is another aspect of our physiology that may be relevant to the costs and benefits of different behaviors. Throughout evolutionary history, being sick could be detrimental to fitness. At best, sickness often means that it is harder to carry out tasks required for survival. At worst, sicknesses can lead to death or otherwise prevent future reproduction (e.g., by causing infertility). If you get sick easily (e.g., example because your immune system does not fight off diseases well), then being overly sociable or sexually active can lead to more costs than benefits. But if you don't get sick easily (e.g., because your immune system is strong) then sociable or sexually promiscuous strategies can lead to benefits in the form of abundant relationship ties or mating success.

In the language of costs and benefits, greater health can make it easier to reap the benefits and avoid the costs of extraverted behavioral strategies. There is little research directly examining relationships between health and personality, but some initial research is promising. For example, Mengelkoch et al. (2022) measured markers of immune system functioning, such as white blood cell count and inflammatory response, in a sample of around 400 people. They found that people who were estimated to have higher immune functioning tended to score somewhat higher on Extraversion.

This research on facultative calibration is still in its early stages, and its likelihood is still being debated (e.g., Zietsch, 2016). But it appears to provide some insight into how our physiology might come to be linked to personality.

Importantly, facultative calibration provides a clear framework for making predictions based on an analysis of the costs and benefits of different personality traits. If any of the physical traits are influenced by genetic differences (which most are), then these differences may ultimately account for some of the genetic variance in personality. Further, if the psychological systems that calibrate personality are assessing the costs and benefits of traits throughout development, then this personality development mechanism can also account for the environmental components of personality variance as well. Ultimately, the facultative calibration hypothesis of personality variation is about more than just physical features like attractiveness, formidability, or health. Any characteristics of a person or their environment can alter the costs and benefits of different personality traits.

Conclusions About Physiology and Personality

The research we have discussed in this chapter makes one point very clear: the links between our physiology and personality should not be expected to be simple 1-1 relationships. Personality traits are complex phenomena that are not simply reflections of physiological differences. In fact, because personality traits as they are currently studied are rooted in language terms for individual differences, there is no guarantee that they will strongly map onto any underlying biological reality. We may need to reexamine how we define and measure personality traits to make further progress in understanding their physiological underpinnings.

Acknowledgements: Immense thanks to my colleague Dr. Chris Miller, a neuroscientist at Fresno State, whose insightful comments and additions greatly improved this chapter. Dr. Miller’s “NOMAD” lab studies the neuroscience of mood and anxiety-related disorders. Check out the NOMAD lab website for more information: www.nomadresearchlab.org/

References

- Cai, H., Zhu, J., & Yu, Y. (2020). Robust prediction of individual personality from brain functional connectome. *Social cognitive and affective neuroscience*, 15(3), 359-369.
- Chen, Y. W., & Canli, T. (2022). “Nothing to see here”: No structural brain differences as a function of the Big Five personality traits from a systematic review and meta-analysis. *Personality Neuroscience*, 5, e8.
- DeYoung, C. G., Beaty, R. E., Genç, E., Latzman, R. D., Passamonti, L., Servaas, M. N., ... & Wacker, J. (2022). Personality neuroscience: An emerging field with bright prospects. *Personality science*, 3.

- Durkee, P. K., Goetz, A. T., & Lukaszewski, A. W. (2018). Formidability assessment mechanisms: Examining their speed and automaticity. *Evolution and Human Behavior*, 39(2), 170-178.
- Geniole, S. N., Bird, B. M., McVittie, J. S., Purcell, R. B., Archer, J., & Carré, J. M. (2020). Is testosterone linked to human aggression? A meta-analytic examination of the relationship between baseline, dynamic, and manipulated testosterone on human aggression. *Hormones and behavior*, 123, 104644.
- Gray, J. C., Owens, M. M., Hyatt, C. S., & Miller, J. D. (2018). No evidence for morphometric associations of the amygdala and hippocampus with the five-factor model personality traits in relatively healthy young adults. *PloS one*, 13(9), e0204011.
- Grotzinger, A. D., Briley, D. A., Engelhardt, L. E., Mann, F. D., Patterson, M. W., Tackett, J. L., ... & Harden, K. P. (2018). Genetic and environmental influences on pubertal hormones in human hair across development. *Psychoneuroendocrinology*, 90, 76-84.
- Herculano-Houzel, S. (2012). The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proceedings of the National Academy of Sciences*, 109, 10661-10668.
- Hyatt, C. S., Sharpe, B. M., Owens, M. M., Listyg, B. S., Carter, N. T., Lynam, D. R., & Miller, J. D. (2022). Searching high and low for meaningful and replicable morphometric correlates of personality. *Journal of Personality and Social Psychology*, 123(2), 463.
- Hyatt, C. S., Hallowell, E. S., Owens, M. M., Weiss, B. M., Sweet, L. H., & Miller, J. D. (2020). An fMRI investigation of the relations between Extraversion, internalizing psychopathology, and neural activation following reward receipt in the Human Connectome Project sample. *Personality Neuroscience*, 3, e13.
- Lin, J., Li, L., Pan, N., Liu, X., Zhang, X., Suo, X., ... & Gong, Q. (2023). Neural correlates of neuroticism: a coordinate-based meta-analysis of resting-state functional brain imaging studies. *Neuroscience & Biobehavioral Reviews*, 105055.
- Manson, J. H., Chua, K. J., Rodriguez, N. N., Barlev, M., Durkee, P. K., & Lukaszewski, A. W. (2023). Sex differences in fearful personality traits are mediated by physical strength. *Social Psychological and Personality Science*, 14(6), 763-771.
- Mengelkoch, S., Gassen, J., Corrigan, E. K., & Hill, S. E. (2022). Exploring the links between personality and immune function. *Personality and individual differences*, 184, 111179.
- Noordhof, A., Kamphuis, J. H., Sellbom, M., Eigenhuis, A., & Bagby, R. M. (2018). Change in self-reported personality during major depressive disorder treatment: A reanalysis of treatment studies from a demoralization perspective. *Personality Disorders: Theory, Research, and Treatment*, 9(1), 93-100. <https://doi.org/10.1037/per0000238>

Rietschel, L., Streit, F., Zhu, G., McAloney, K., Frank, J., Couvy-Duchesne, B., ... & Rietschel, M. (2017). Hair cortisol in twins: Heritability and genetic overlap with psychological variables and stress-system genes. *Scientific reports*, 7(1), 15351.

Stein, D. J., & Stahl, S. (2000). Serotonin and anxiety: current models. *International clinical psychopharmacology*, 15, S1-S6.

Sundin, Z. W., Chopik, W. J., Welker, K. M., Ascigil, E., Brandes, C. M., Chin, K., ... & Tackett, J. L. (2021). Estimating the associations between big five personality traits, testosterone, and cortisol. *Adaptive Human Behavior and Physiology*, 1-34.

Sugiyama, L. S. (2015). Physical attractiveness in adaptationist perspective. *The handbook of evolutionary psychology*, 292-343.

Zietsch, B. P. (2016). Individual differences as the output of evolved calibration mechanisms: does the theory make sense in view of empirical observations? *Current Opinion in Psychology*, 7, 71-75.

9 Emotions and Personality

Personality traits as we commonly measure them in research are descriptive summaries of ways that people tend to differ from one another in their tendencies of thought, feeling, and behavior. But what generates these thoughts, feelings, and behavior? We can't just say that personality traits generate these differences because that would be circular reasoning, which is a [logical fallacy](#). To explain personality, we ultimately need to tie these abstract trait constructs back to some more basic entities in the human mind. In this chapter, we'll look at a set of entities that are candidates for helping us better understand how personality variation may be generated: emotions.

What are emotions? You undoubtedly have some experience with emotions, and you may have your own views about how they work. In the field of psychology there are many theoretical views about what emotions are—how many there are; how they work; whether they have positive or negative “valence”; whether they feel good or bad or hot or cold. Much of this debate relies too much on the subjective experience and folk descriptions of emotions. For our purposes in this chapter, we'll take a particular computational view of emotions that is informed by evolutionary psychology, which provides a structured and non-arbitrary way of identifying and understanding emotions.

Note: If phrases like “adaptations”, “adaptive problems”, “evolutionary fitness” are unfamiliar to you, you may want to review the chapter on “Evolution and Personality” before continuing.

An Evolutionary Framework for Studying Emotions

Emotions can be viewed as adaptations that coordinate a wide array of cognitive and physiological processes within organisms in response to specific adaptive problems (Cosmides & Tooby, 2000). Under this perspective, the features of many emotions are expected to have been designed through the iterative tinkering processes of natural selection and sexual selection, so that their features would *on average*¹ produce behavior that helps to solve adaptive problems, and ultimately, increase evolutionary fitness. Different emotions are expected to have evolved in response to different adaptive challenges, each producing different adaptive responses. After all, what counts as fitness-enhancing behavior differs depending on the adaptive

¹This means that emotions should not be expected to produce the perfectly adaptive behavior in every situation or every person, even when they are functioning just as they are “supposed” to. Instead, the responses that are produced can be expected to be those that have worked most frequently across most situations throughout evolutionary history.

problem being faced — hugging an approaching predator almost certainly isn't an adaptive response, but hugging an approaching loved one probably often is.




Figure 9.1: Emotions as a particular class of psychological mechanisms.

Using this view of emotions, we can think about emotions as having an input-output structure, like the one shown in Figure 9.1. What we call emotions are the patterns of activation that reliably occur in response to inputs that would have been reliably associated with the presence of adaptive problems in the ancestral past. Based on these inputs, emotions assess the severity of the adaptive problem and produce outputs—like thoughts, feelings, and behaviors—that would have reliably helped to solve the adaptive problem, on average, over evolutionary history. Table 9.1 provides examples of different adaptive problems.

Table 9.1: Examples of adaptive problems along with their potential cues and solutions

Adaptive Problem	Cues	Solutions
Depleted nutritional state	Low blood sugar, empty stomach	Find food, eat
Predator threats	Fangs, large paw prints, growling	Escape, fight, heightened caution
Devaluation by conspecifics	Avoidance, harm, not sharing	Attempt to regain value
Losing valued cooperation partners	Death of cooperator, mate left	Attempt to regain partner, seek new partner

Let's explore the adaptive problem of predator threats in more detail. The inputs to an emotion that deals with predator threats should be activated by information from the external environment (inputs) that are reliable clues to the presence of predators. There are many

potential clues that can range in their severity or information value for assessing whether there is a predator threat present. For example, seeing large paw prints or scat from a predator might raise the probability of a predator threat a little bit—but seeing the predator’s fangs and hearing it growl is much a more informative cue and raises the probability of a predator threat a lot! So, the inputs that this predator threat emotion system allow it to assess the level or severity of the predator threat that is present.

In response to this assessment, the predator threat emotion system produces outputs in the form of coordinated responses that could help to mitigate the threat. In theory these responses should be proportional to the level of threat (input) detected because having too high a response could waste energy or lead to other negative outcomes, and having too low a response could lead to the threat not being averted—both of which are bad for fitness. For example, if scat or paw prints are all that is around and the emotion system assesses the threat levels as “medium”, then a proportional response might be to freeze and listen or leave the area or have heightened caution going forward. But if fangs and growling are present and the emotion system assesses the threat level as “very high”, then a proportional response might be trying to flee or call for help or fight.

We might call this predator threat emotion system “fear”. But it’s important to recognize that the names of emotions that exist in any given language may not necessarily correspond to evolutionary reality of emotions—words are just useful tools to describe concepts. What we call “fear” in language could refer to several different emotional systems within the brain, each designed to respond to a particular adaptive problem. Some of what we call fear, might be referring to the activity of a specific psychological mechanism (emotion) that is designed to respond to predator threats in particular. But other experiences that we call “fear” could be referring to the activity of emotion systems that are more narrowly dedicated to adaptive problems such as threats from venomous snakes, poisonous insects, physically dangerous conspecifics, and cliffs or other high-up places. We can make more detailed descriptions and boundaries between emotions that are less arbitrary than those that our language and folk psychology makes by starting with adaptive problems, their inputs, and their outputs.

We could follow similar logic for delineating the potential inputs and outputs for any adaptive problem to develop descriptions of other potential emotions. For example, if the adaptive problem is [maintaining enough nutrients to power the body], then an input could be [low blood sugar] and potential outputs would be [stomach rumbling], [find food], or [eat food]. We might call the subjective experience of the activity of this emotion “hunger” (Al-Shawaf et al., 2016).

Emotions Coordinate Many Processes

Emotions are expected to coordinate a wide array of processes in response to adaptive problems, including those that are beyond what we might typically think of as behavioral outputs. Some of these processes are listed in Figure 9.2. For example, the activation of predator fear

emotion system could change our current goals and motivations from finding food to escaping an area. Further, activation could shift our attention towards cues of predator presence and trigger information-gathering to further assess the threat level. Moreover, the concepts we apply to objects in the environment could change—what was just a rock before we detected a predator threat might now be conceptualized as a good weapon. Perception might change such that time feels slowed down or recall memory might be momentarily enhanced and it is easier to remember escape routes. Physiological changes, such as increased heart rate to pump oxygen to muscles in preparation for action, are also likely a component of many emotional responses. Many human emotions might also coordinate outputs that allow for communication of relevant information, via verbal communication or even non-verbal facial expressions; for example, fear might involve verbally communicating to others about the danger or making a face that communicates the information nonverbally.




Figure 9.2: Examples of responses that could be coordinated by emotions.

Importantly, not all emotions should be expected to trigger all of these processes in all instances. The responses that are coordinated by a given emotion will depend on the particular set of responses that would have been necessary to solve the problem on average throughout the emotion’s evolutionary history. For example, if nonverbal communication is not an important or detrimental for addressing a particular adaptive problem, then we should not expect there to be a clear facial expression associated with the emotion.

Emotions Have a Species Typical Design

In general, this evolutionary framework of emotions expects that emotion systems function in basically the same way across all humans—that is, there is a species typical design of each emotion. This means that the ways emotions work should be relatively the same across all cultures. But of course, the words and descriptions that cultures develop to label, describe, and explain can differ substantially. This is because all humans who have ever lived share common ancestry in Africa around 70,000 years ago, which is very recent in evolutionary time. There likely hasn’t been enough time for different emotions—which are complex adaptations

requiring many pieces of our cognition and physiology to work together in concert—to evolve in the short time since humans have migrated to geographically and genetically isolated parts of the world, especially because there is still substantial gene flow between populations.

Researchers have applied this framework of emotions to study a wide array of emotions, including anger (Sell, 2011), shame and pride (Sznycer et al., 2016; 2017), jealousy (Buss & Haselton, 2005), fear (Öhman & Mineka, 2001), lassitude (the emotion of being sick; Schrock et al., 2020), gratitude (McCullough et al., 2008), guilt (Gilbert, 2003), disgust (Tybur et al., 2013), to name a few. I encourage you to at least skim one of the papers cited above for an emotion you’re curious about to see how this approach works in practice. Next, we’ll use this emotion framework to attempt to understand individual differences and personality.

Individual Differences in Emotions

You might be confused about how emotions could have a universal or species-typical design and yet also be expected to be useful for explaining individual differences. But just like humans generally have the same species-typical body structure without looking exactly the same, emotions might have the same structure and still have minor differences. So, there is lots of room for individual differences in emotions! We’ll explore below how individual differences in emotions and emotional outputs might arise and become relevant to personality.

Many Moving Parts

Emotions depend on many interactions between many cognitive systems. For example, each emotion system may rely on “detecting algorithms” that monitor the situations for cues of adaptive problems and combine the cues to develop probabilistic estimates. There could well be subtle individual differences in the functioning or accuracy of these systems leading to small (or maybe even large) differences in assessments of adaptive problem levels. Other differences could be in the sub-systems that handle prioritizing goals, or the algorithms that calculate appropriate behavioral or physiological responses. Because differences in any of these subsystems within an emotion will play a role in determining behavior, small differences in any systems can lead to observable differences in behavior.

Many Sources of Variation

As with any trait, differences in emotion systems can arise through genetics or environmental influences. Small genetic mutations can lead to random differences or “noise” in the functioning of emotion systems that may increase or decrease their effectiveness. Environmental insults like radiation or brain damage could have similar detrimental effects on emotion systems. Environmental influences across development could also lead to differences if emotion systems

are calibrated (fine-tuned) by environmental information. For example, if the anger system gets feedback through development that anger is an effective bargaining tool, the anger system might be more easily triggered; or if the anger system gets feedback through development that anger is not an effective bargaining tool, then it may be less easily triggered. In sum, there are many ways that individual differences in emotion systems can arise, but this is not exactly the focus of the current module so let's move on.

The Social Relevance of Differences in Emotions

Regardless of how the differences in emotion systems arise, the differences can be expected to result in differences in patterns of thoughts, feelings, and behavior (i.e., personality). Some people will experience emotions (e.g., anger, pride, shame, gratitude, guilt, fear) and their associated outputs (e.g., behavior, feelings) more easily or more often than others. These individual differences are useful for others and ourselves to recognize and label because this information may help us predict *future* behavior, which could allow us to make better decisions about social interactions. For example, if you know someone gets really angry really easily, you know to avoid doing anything to make them angry. Or if you know someone gets scared very easily, they may not be the best person to invite on a dangerous expedition.

Thus, the words that have emerged in the language may effectively capture socially relevant individual differences in emotions. Because words are the basic building blocks of the abstract personality descriptions that are popular in contemporary personality research² (e.g., Big Five and HEXACO), personality trait constructs may ultimately capture socially relevant variation in emotion systems.

We should therefore expect that individual differences in emotion systems and individual differences in personality constructs should share a lot of overlap. Armed with this appreciation, we can start to develop predictions about how variation in emotion systems might relate to behavior, social judgments, and ultimately broad personality traits like those in the Big Five or HEXACO models.

Reverse Engineering Personality Constructs via Emotions

If variation in emotions is roughly captured in person-descriptor words, and these person-descriptor words led to personality constructs (c.f., the lexical approach), we should be able to reverse engineer personality constructs (e.g., neuroticism, agreeableness) by linking them back to differences in emotion systems. Evolutionary psychologists (including me) have worked out a framework for tackling this endeavor (Lukaszewski, 2020; Lukaszewski et al., 2020). I'll summarize two strategies for reverse-engineering personality constructs here.

²See the lexical studies discussed in the "Contemporary Personality Perspectives" lecture for a refresher.

Bottom-Up: From Emotions to Personality Constructs

The first strategy, which could be called “bottom-up”, starts individual differences in emotions and then links them to established personality trait constructs. Here, the first step would be to identify an adaptive problem, develop hypotheses about the design features of an emotion system that could solve that problem, and test those predictions to provide evidence for the hypothesized structure of the emotion. Either in these initial tests or in follow-up studies researchers can explicitly estimate individual differences in the emotion system, such as differences in the reactivity of the emotional response, or severity of the emotional response. The researcher can make predictions about how these individual differences in the emotional responses and manifest behavior may relate to social judgments and person-descriptor terms. These predictions can then be tested to see if the measured differences in emotions predict personality traits.

Top-Down: From Personality Constructs to Emotions

The second strategy, which could be called “top-down”, starts with established personality constructs and works backwards to link them to individual differences in emotions. Here, the first step would be choosing the focal personality trait(s) for study. Then, researchers can make predictions about what sorts of emotional reactions the personality trait(s) constructs may be capturing. Then the researchers can collect data on individual differences in emotions and personality traits to see if they are linked in the predicted ways. You should be able to practice this approach thinking about any trait you might be interested in—go ahead and give it a try!

Empirical Examples

Researchers have begun to apply this framework linking emotions to personality. Let’s look at a few empirical examples here.

Anger and Agreeableness

The anger system is hypothesized to be designed to bargain for better treatment from people when they undervalue you (Sell, 2011). In line with this, anger is activated by information that suggests a person undervalues you (e.g., they harm you to get something they want). And the anger system produces outputs that effectively bargain for better treatment in the future, such as aggressive threats, demanding better treatment, or threatening to withdraw benefits in a relationship. In a sample of almost 400 people, self-reported individual differences in anger proneness correlate with self-reported individual differences in agreeableness at 0.79 (Lukaszewski et al., 2020). This suggests that over half of variation in the personality construct of agreeableness can be linked back to individual differences in the anger system.

Jealousy and Neuroticism

The jealousy system is hypothesized to be designed to prevent or mitigate threats to a valued relationship (e.g., romantic relationship or other cooperative relationship). Jealousy is activated by cues that a relationship may be threatened, such as your partner being interested in someone else or others being interested in your partner. The outputs of jealousy seem to be designed to effectively stop the relationship threats, for example by “guarding” your partner or increasing investment in a partner so that they are less likely to seek investment elsewhere. In a currently unpublished study, we examined jealousy in a sample of around 200 people in romantic relationships. We found that the severity of a person’s jealousy reaction in response to imagining their partner engaging in behaviors with different levels of potential infidelity risk was correlated with their own self-perceptions of Neuroticism as well as others’ perceptions.

Emotions Provide Insight into Lexical Personality Constructs

Although this research is still in its early stages, there is promising evidence that examining emotional variation will be useful for developing better descriptions and explanations of lexical personality constructs. My lab is currently exploring how individual differences in emotions like guilt, pride, shame, gratitude, and disgust relate to the personality constructs in the HEXACO framework. This will allow us to link the complex phenomenology of personality back to more basic constructs that we already have a relatively good understanding of at the *computational* and *algorithmic* levels of analysis. In the future, this may make it easier to develop better predictions about the *implementational* level of analysis for personality traits.

By adding up all the small differences in emotion systems, we may be able to account for large amounts of differences in broader personality constructs. This approach again shows that it is often necessary—or at least useful—to combine both ultimate and proximate perspectives to understand psychological phenomena. An understanding of the ultimate, evolved functions of emotions allows us to better examine how proximal differences in the functioning of emotions arise and how they relate to personality.

Of course, personality variation is probably not *only* emotion variation. There are likely many other psychological entities that generate differences in thoughts, feelings, and behavior. Other such entities in the mind might include things like inferences, preferences, values, signal detection thresholds, decision rules, and motivational systems. We’ll explore some of these in future modules.

References

Al-Shawaf, L. (2016). The evolutionary psychology of hunger. *Appetite*, 105, 591-595.

- Buss, D. M., & Haselton, M. (2005). The evolution of jealousy. *Trends in Cognitive Sciences*, 9(11), 506-506.
- Cosmides, L., & Tooby, J. (2000). Evolutionary psychology and the emotions. *Handbook of emotions*, 2(2), 91-115.
- Gilbert, P. (2003). Evolution, social roles, and the differences in shame and guilt. *Social Research: An International Quarterly*, 70(4), 1205-1230.
- Lukaszewski, A. W. (2020). Evolutionary perspectives on the mechanistic underpinnings of personality. In J. Rauthmann (Ed.), *The handbook of personality dynamics and processes* (pp. 523-550). San Diego, CA: Elsevier Press
- Lukaszewski, A. W., Lewis, D. M., Durkee, P. K., Sell, A. N., Sznycer, D., & Buss, D. M. (2020). An adaptationist framework for personality science. *European Journal of Personality*, 34(6), 1151-1174.
- McCullough, M. E., Kimeldorf, M. B., & Cohen, A. D. (2008). An adaptation for altruism: The social causes, social effects, and social evolution of gratitude. *Current directions in psychological science*, 17(4), 281-285.
- Öhman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: toward an evolved module of fear and fear learning. *Psychological review*, 108(3), 483.
- Sell, A. N. (2011). The recalibrational theory and violent anger. *Aggression and violent behavior*, 16(5), 381-389.
- Schrock, J. M., Snodgrass, J. J., & Sugiyama, L. S. (2020). Lassitude: The emotion of being sick. *Evolution and Human Behavior*, 41(1), 44-57.
- Sznycer, D., Al-Shawaf, L., Bereby-Meyer, Y., Curry, O. S., De Smet, D., Ermer, E., ... & Tooby, J. (2017). Cross-cultural regularities in the cognitive architecture of pride. *Proceedings of the National Academy of Sciences*, 114(8), 1874-1879.
- Sznycer, D., Tooby, J., Cosmides, L., Porat, R., Shalvi, S., & Halperin, E. (2016). Shame closely tracks the threat of devaluation by others, even across cultures. *Proceedings of the National Academy of Sciences*, 113(10), 2625-2630.
- Tybur, J. M., Lieberman, D., Kurzban, R., & DeScioli, P. (2013). Disgust: evolved function and structure. *Psychological review*, 120(1), 65.

10 Motivation and Personality

In the last module, we explored how variation in personality maps on to variation in emotions. We learned that ultimate perspectives about the evolved functions of emotions can inform proximal research questions about relationships between emotional variation and personality. In this chapter, we'll use another ultimate framework—that of evolved fundamental motives—to explore how differences in motivational priorities might help to explain abstract personality constructs captured by descriptive lexical dimensions like the Big Five and HEXACO traits.

Note: If phrases like “adaptations”, “adaptive problems”, “evolutionary fitness” are unfamiliar to you, you may want to review the chapter on “Evolution and Personality” before continuing.

What Are Motivations?

When I ask students what a motivation is, I often get answers that involve “wanting”, “desiring”, “needing”, or “striving” for a goal. This isn't wrong. It certainly fits with my own experience of what motivations feel like—and probably with your subjective experience as well. But we want to delve a little deeper to try to get beyond the subjective experience of motivation and the words we use to describe these experiences. We need a more *mechanistic* view of motivations: one that treats motivation as a component of the broader computational system that makes up the human mind.

For our purposes, motivations are internal states that activate and direct behavior towards goals or needs. Under this view, motivations aren't necessarily consciously represented. In fact, they probably aren't most of the time. Most of our motivations are totally outside of our awareness, implicitly guiding behavior towards goals and needs. We'll talk more about what these goals and needs might be later. But what are these internal states and how might we think about them?




Figure 10.1: Diagram of pulley system. (public domain [image hyperlink](#))

Cognitive psychologists in the evolutionary psychology tradition developed the concept of **Internal Regulatory Variables**, which provide a more detailed way of thinking about what motivations are from a computational cognitive science perspective (Tooby et al., 2008). An internal regulatory variable (abbreviated as IRV) can be thought of as a numerical representation of value or meaning about states of the world. You might want to think about IRVs like parameters in a mathematical equation or computer code; or as rocks in some sort of old-timey physical pulley system like that shown in Figure 10.1. Depending on the magnitude of the parameters or the weight of the rocks, they lead to different consequences for the equation, computer program, or pulley system.

Internal regulatory variables could represent many types of information. Examples include: your own physical strength or attractiveness; the physical strength, attractiveness, or genetic relatedness of conspecifics; how much a person values or disvalues us; how long it has been since you heard a scary noise in the bushes; the fairness of resource distributions within a group; the nutritional value of an edible object; status impacts of different behaviors; and many more features of the world. If IRVs exist within the mind, they can be involved in regulating behaviors, thoughts, and feelings.

For example, imagine a situation where a person has just taken an apple from you. The decision of whether to take action to get apple back, and which actions to take, depend on many other details of the situation: who is the person to you—friend or foe? Do you need the apple for sustenance or is it superfluous? If the apple is needed, are you able to physically

take the apple back from the person? How much do others value the person? Do you have kin nearby who would back you up in a conflict? How would losing such a conflict impact your status within the community? Each of these aspects of the situation (and many others) can be represented by IRVs in the mind that are involved in the computation of decisions to approach the person, bargain with the person, fight the person, etc. — which we experience and lexically label with our everyday language as “desire” or “motivation”.

Internal regulatory variables play a role in regulating behavior, thoughts, and feelings in concert with other entities in the mind, like emotions that we discussed in the previous chapter. The folk concept of motivation and the representations of it in our language may reflect individual differences in underlying internal regulatory variables. I doubt that all IRV’s are reflected in what we call motivations—but it seems likely to me that all motivations are manifestations of IRVs. We experience (some of) the computations carried out by our minds subjectively as “wanting” or “needing” or “being motivated” to do things (e.g., do homework, brush our teeth, work out, go on a date, hang out with friends, write a free personality textbook).

The theoretical concept of motivation that we’ll rely on is the different weights that our minds place on the value of different goals. But what are the goals that humans have? How can we determine which ones are important to study?

What Goals Are Motivations For?

There have been many attempts to identify the most basic or important human goals or motivations. A common theme in motivation theories is that peoples’ behavior is largely determined by the goals they are pursuing and the motives that are active.

Early Motivation Theorists

Freud and the neo-Freudians that we discussed back in the “Past Personality Perspectives” chapter each had their own hunches about which motives drive for human behavior. Freud thought that the libido, or the sex drive, was the primary motivational drives in humans. As for the neo-Freudians, Karen Horney argued that the primary human drives were the need for belonging and security, while Alfred Adler argued that people were primarily motivated by the need for superiority (or at least avoiding inferiority). Who is right?

These theories of motivation appear to overlook the fact that there are likely to be many fundamental human motives. There probably isn’t just one motive or drive that can be held responsible for all of the varied behaviors that people engage in. We probably need several motives to explain the many in thoughts, feelings, and behaviors that humans produce.

Maslow's Hierarchy of Needs

One of the first proposals of a list of multiple fundamental motives was developed by Abraham Maslow in 1943, which is referred to as Maslow's Hierarchy of Needs¹. Figure 10.2 shows the typical representation of Maslow's hierarchy as a pyramid. According to Maslow, behavior will be geared towards addressing whichever goal is currently being worked on, and the lower-order needs at the bottom of the hierarchy must be fulfilled before the higher-order needs (higher on the pyramid) can be achieved.




Figure 10.2: Maslow's Hierarchy of Needs depicted as a pyramid, with lower-level primary needs at the bottom and higher-order secondary needs towards to the top.

Physiological needs are basic needs of survival, such as water, food, shelter, and rest. Moving up the pyramid, the next level is safety needs, which involve seeking protection from physical and emotional harm, as well as security and stability. Beyond this, the third level is composed of social needs, highlighting the importance of belonging, companionship, and meaningful relationships. Esteem needs at the fourth level, revolve around self-esteem, self-respect, and recognition from others. Finally, at the pinnacle of the hierarchy, we find self-actualization needs, representing things like aspiration for personal growth, spiritual enlightenment, self-fulfillment, and the realization of one's ideal self.

¹Note that psychologists use "goals", "motivations", and "needs" pretty much interchangeably.

Maslow's hierarchy was primarily about understanding and explaining differences in fulfillment, happiness, and well-being. But we can apply it to personality as well. If a person's behaviors, thoughts, and feelings are dominated by the need to fulfill their current need, then their personality could be explained by looking at which needs the person is typically confronted with addressing. A person who is often simply trying to fulfill physiological needs can be expected to act different, on average, from a person who has met their physiological needs and is striving for higher needs, like self-actualization. For example, we might expect a person who does not have their basic safety needs met, may act more worried or impatient—ultimately being perceived as more neurotic or less agreeable. A person who has all their basic needs met and is working on self-actualization needs might engage more with complex, abstract ideas or be more willing to donate their time—leading them to be perceived as open and agreeable.

Maslow's hierarchy is certainly a popular framework of motivation and it can be useful in understanding behavior and individual differences, in theory. But the needs in the Maslow's hierarchy are ultimately somewhat arbitrary. Maslow didn't arrive at the particular needs in his hierarchy through any broader theory and he didn't identify the particular needs in the hierarchy through systematic data collection. Ideally, we would want a theory-driven framework that explains why humans have the motives they do and that is supported by empirical data. Next, we'll explore just such a framework and see how it can help us understand personality variation.

The Revised Fundamental Motives Framework

In the early 2000s evolutionary psychologists revised Maslow's Hierarchy based on evolutionary theory and considerations of fundamental adaptive problems (Kenrick et al., 2010). If complex motivations—sets of internal regulatory variables—evolve in humans, it can be expected that they helped solve some problem of survival and reproduction over the course of evolutionary history. Thus, the set of human goals and motivations should largely correspond to the set of adaptive problems that humans faced throughout our evolutionary history. Based on considerations of adaptive problems, Kenrick and his colleagues identified the fundamental motives depicted in Figure 10.3.




Figure 10.3: Revised fundamental motives based on consideration of adaptive problems faced throughout human evolution.

Under this evolutionary fundamental motives framework, a given motive will be activated in response to the relevant cues of the adaptive problem being present. For example, if you enter into a competition that could affect your status, then internal regulatory variables related to status motives will be relevant; or if predators are detected, then internal regulatory variables related to self-protection will be involved. Of course, some goals, like kin care or mate acquisition and retention can only be activated at particular life-stages and depend on achieving some other goals—mate retention is only relevant if you already solved the adaptive problem of acquiring a mate, for example. Moreover, these motives are fairly rough categories that could be further divided up into more narrowly defined adaptive problems. But ultimately, this framework provides some useful dimensions that we can explore in relation to personality.

Linking Fundamental Motives to Personality Constructs

Neel and colleagues (2016) developed scales to measure individual differences in each of the fundamental motives, along with more specific aspects of them (e.g., disease avoidance within safety needs). Evidence suggests these scales have good reliability and their validity has largely been established by linking them to other individual difference constructs like personality traits. As shown in Table 10.1, several studies in sample from the USA ($Ns > 200$) show that

individual differences in fundamental motives are reliably linked to Big Five personality trait constructs in predictable ways (Neel et al., 2016).

Table 10.1: Correlations between fundamental motives and Big Five personality traits from Neel et al. (2016)

Fundamental motive	E	A	C	N	O
Self-Protection	.04	.03	.09*	.17*	.06*
Disease Avoidance	-.07*	-.10*	.10*	.10*	-.03
Affiliation (Group)	.41*	.45*	.21*	-.22*	.14*
Affiliation (Exclusion Concern)	-.14*	-.14*	-.25*	.43*	-.08*
Status	.26*	-.06*	.03	.06	.12*
Mate Seeking	-.01	-.16*	-.13*	.05	.03
Mate Retention (General)	.03	.29*	.29*	-.09*	.12*
Mate Retention (Breakup Concern)	-.13*	-.25*	-.31*	.37*	-.03
Kin Care (Family)	.21*	.42*	.29*	-.15*	.11*
Kin Care (Child)	.02	.23*	.19*	.03	.08

Limitations of Current Motive Measures

Of course, this research can't really show that motivations are *causing* the variation in personality. It could also be that differences in personality lead people to experience different situations that have more of particular adaptive problems. For example, extraverted individuals are more likely to be in social situations, so they may be more likely to experience cues to adaptive problems relevant to status and affiliation. Or other individual differences like age may cause differences in both the fundamental motives and personality traits, leading them to be correlated even though they do not have a direct influence on one another.

Moreover, because both personality questionnaires and the fundamental motives questionnaires rely on self-reports in response to simple questions, it is possible that some of the correlations between personality and motivations are just reflective of having similar questions for each measurement tool. As an exercise in evaluating psychological measures, compare the example questions in Table 10.2 to questions in popular Big Five personality inventories—do you think there is a lot of overlap in the language of the questions that could drive the observed correlations in the data?

Table 10.2: Example questions from the fundamental motives questionnaire. Items from Neel et al. (2016).

Subscale	Sample item
Self-Protection	I think a lot about how to stay safe from dangerous people.
Disease Avoidance	I avoid people who might have a contagious illness.
Affiliation (Group)	I like being part of a team.
Affiliation (Exclusion Concern)	I would be extremely hurt if a friend excluded me.
Status Seeking	I want to be in a position of leadership.
Mate Seeking	I am interested in finding a new romantic/sexual partner.
Mate Retention	It is important to me that my partner is emotionally loyal to me.
Breakup Concern	I often think about whether my partner will leave me.
Kin Care (Family)	Caring for family members is important to me.
Kin Care (Children)	I like to spend time with my children.

In my own research (Durkee et al., 2023), I’ve also used self-report scales to measure HEXACO traits and status motivations in samples of participants in India and the USA ($Ns > 800$). I found that higher status motivations were reliably connected to higher Extraversion and lower Honesty-Humility in both countries. In India, higher status motivations were also linked to higher Conscientiousness and higher Openness. These findings add some additional evidence that motivational priorities are reliably linked to personality traits. But they also suggest that some of these connections may differ depending on the cultural interpretations of motivations and their manifest behavior, thoughts, and feelings.

Closing Thoughts About Motivations and Personality

In this chapter, we've learned that the concept of motivations may refer largely to subtle differences in the weight that organisms—or more accurately, organisms' internal regulatory variables—place on different adaptive problems. We explored how differences in motivations could be expected to relate to descriptive personality constructs like the Big Five and HEX-ACO traits. Just like with emotions, many of the personality trait terms we employ describe subtle differences in peoples' motivational priorities, or the weights that their mind places on addressing different adaptive problems. This can again help us to better predict how others may behave in different situations.

Ultimately, we will need more direct measures of the state of internal regulatory variables to examine their roles in personality variation more accurately. But even at the current level of analysis that relies on self-reports, we can get a sense of how motivation and personality constructs are intertwined. This gives us a better understanding of both constructs.

The motivational frameworks we explored here are also just useful tools for understanding individual differences in thoughts, feelings, and behaviors—including our own! Next time you feel (un)motivated to do something (e.g., studying), you can think about what how the situation may relate to a fundamental adaptive problem and imagine different internal regulatory variables interacting in your mind to produce your experience. Maybe you can identify characteristics of the situation or yourself that you can make more salient or change in order to alter your motivations to produce behaviors, thoughts, and feelings, that you think are desirable. As we'll see in a later chapter on personality change, if you can repeatedly trigger and act on specific motivations over long periods of time, you may be able eventually change the average trajectories of your thoughts, feelings, and behavior.

References

- Durkee, P. K., Lukaszewski, A. W., & Buss, D. M. (2023). Status-impact assessment: Is accuracy linked with status motivations?. *Evolutionary Human Sciences*, 5, e17.
- Kenrick, D. T., Griskevicius, V., Neuberg, S. L., & Schaller, M. (2010). Renovating the pyramid of needs: Contemporary extensions built upon ancient foundations. *Perspectives on psychological science*, 5(3), 292-314.
- Maslow A.H. (1943). A theory of human motivation. *Psychological Review*, 50, 370–396.
- Neel, R., Kenrick, D. T., White, A. E., & Neuberg, S. L. (2016). Individual differences in fundamental social motives. *Journal of personality and social psychology*, 110(6), 887.
- Tooby, J., Cosmides, L., Sell, A., Lieberman, D., & Sznycer, D. (2008). Internal regulatory variables and the design of human motivation: A computational and evolutionary approach. *Handbook of approach and avoidance motivation*, 15, 251.

11 Development and Personality

This chapter delves into the topic of personality development from infancy to adulthood. We'll examine personality development through the lens of the Big Five personality traits. Specifically, we'll explore how the structure of personality changes across development, how personality becomes more stable over time, and normative trends in how personality “matures” as we get older.

Temperament versus Personality

Historically, developmental psychologists have distinguished between *temperament* and personality (Shiner et al., 2021). Temperament in this context refers to individual differences in the first two years of life. In theory, temperament is expected to have a stronger biological basis because these individual differences emerge soon after birth, ostensibly before culture and other aspects of the environment can have much influence.

Distinguishing between temperament and personality as different constructs might not be necessary, however. Shiner and colleagues noted in a recent review of personality development that “it has become increasingly clear that temperament versus personality is a distinction without much of a difference” (Shiner et al., 2021, p. 117).

For our purposes, we will refer to all individual differences in tendencies of thought, feeling, and behavior as personality, regardless of what life stage we are examining. However, as we will see, this doesn't mean that personality looks exactly the same across the lifespan.

The Little Six?

There are some important differences between youth and adult personality (Soto & Tacket, 2015). One of the biggest differences is in the *structure* of personality. We learned about the structure of adult personality back in the “Contemporary Personality Perspectives” module.

To review, personality structure refers to the patterns of overlap among personality traits. The Big Five personality structure, for example, was identified because the many trait terms that describe people can be roughly grouped into five large groups of broad personality traits. This five-factor structure is sufficient to summarize the individual differences between adults

in many cultures. But does the structure of personality look the same in childhood and adolescence?

There are non-trivial differences in the structure of personality in childhood and adulthood. When measuring the Big Five in adults, the dimensions are mostly uncorrelated—meaning they do not overlap very much. In childhood, however, Conscientiousness and Openness exhibit substantial overlap, as do Conscientiousness and Agreeableness. Moreover, a study involving over 3,000 children and early adolescents from five countries found that only extraversion, agreeableness (primarily based on disagreeable behavior), and openness (primarily based on intellectual interests and ability) consistently emerged as broad dimensions of personality across different cultures and age groups (Tacket et al., 2012). This raises doubts about whether the Big Five personality traits adequately capture youth personality. An alternative structure, the *Little Six*, has been suggested as a potential candidate.

The “Little Six” is a model of youth personality that combines key traits from both child temperament and adult personality studies (Soto & John, 2014). Child temperament models have historically included include four dimensions: sociability, negative emotionality, persistence, and activity level. Interestingly, three of these dimensions align well with adult personality traits, like extraversion, neuroticism, and conscientiousness. This suggests that there might be six fundamental personality dimensions: extraversion, agreeableness, conscientiousness, neuroticism, openness to experience, and *activity*. Recent research on parents’ reports of their children’s personalities supported this idea, showing that the Little Six structure is present consistently from middle childhood through adolescence, while the traditional Big Five structure emerges later in late adolescence. These findings also suggest that as we age, childhood variation in the activity dimension merges with Extraversion.

General Principles of Personality Development

There are two broad principles that have been supported by empirical research on personality development. We will call the first the “cumulative continuity” principle. We will call second the “maturity” principle. We will explore each in the relevant sections below, and also examine some interesting caveats.

The cumulative continuity principle is a fundamental concept in the study of personality development, particularly in adults. It refers to the *stability* of personality traits over time in terms of rank-order consistency in personality scores. This is typically examined using test-retest correlations of personality scores across varying periods of time (e.g., 2 weeks, months, or years). Research on the cumulative continuity principle has consistently shown that personality traits tend to become increasingly stable as individuals progress through adulthood (e.g., Bleidorn et al., 2022). This means that individuals tend to maintain their relative positions on various personality dimensions as they age. In simpler terms, if someone is more extraverted than their peers during early adulthood, they are likely to maintain this higher

level of extraversion as they grow older. Figure 11.1 summarizes the cumulative continuity principle.



Figure 11.1: Summary of the cumulative continuity principle.

Personality might stabilize with age for a variety of reasons. Physical features, such as height, tend to stabilize after puberty. If those features play a role in personality, then they can contribute to personality stability. Likewise, peoples' environments tend to stabilize as they get older. For example, people tend to settle into careers that expose them to similar situations that lead to more stability in behavior, thoughts, and feelings.

The idea of *person-environment transactions* provide a useful conceptual tool for thinking about how environments and personality interact to stabilize personality over time (Donnellan, 2023). Let's explore the three types of person-environment transactions and develop an example of each. You should be able to identify examples of these transactions in your own life as well.

Active person-environment transactions capture the fact that people seek out experiences and environments that are consistent with their personality traits. For example, a person who enjoys taking risks will be more likely to seek out risky situations (e.g., casinos) than someone who is risk averse. Over time, this can lead to average differences in environments that contribute to the stability of personality.

Reactive person-environment transactions capture the fact that different people react differently to the same environments because of their personality traits. For example, if you were to drag a group of friends to a loud karaoke bar for your birthday, you may see that some friends have a fun time and enjoy the event, whereas some friends may do all they can to avoid being the center of attention. This tendency for personality traits to be associated with different responses to the same environment also contributes to personality stability.

Evocative person-environment transactions capture the fact that different people may trigger different responses from others in their social environments

because of their traits. For example, a person who is very disagreeable or rude may be more likely to experience others being rude or disagreeable in return, whereas some who is kind and polite may be more likely to elicit kindness and politeness from others. This can create feedback loops between personality and situations that reinforce the average patterns of thoughts, feelings, and behaviors that make up personality.

Caveat: The stability of personality is far from perfect. Research on the stability of personality across the lifespan suggests that the stability of personality increases throughout childhood, our teen years, and early adulthood before plateauing in the early thirties and staying relatively level into older adulthood. Even when the estimates of personality stability are at their highest point, they suggest that about 50% of the variation in the rank-ordering of personality traits at one time can be explained by differences in the rank-ordering of personality in the past. That means that another 50% of the rank-ordering must be attributable to changes in personality (e.g., systematic changes or random changes). Thus, even though our personality stabilizes as we age and reaches a pretty high level of stability, that doesn't mean that changes do not occur.

The maturity principle is another important concept in the field of personality development, particularly in the context of adulthood. It focuses on the changes in mean-level development, or the average levels of specific personality traits at different ages. According to research on personality development, most people across many countries tend to become more agreeable, conscientious, and emotionally stable (less neurotic) as they grow older (Bleidorn, 2015). Figure 11.2 summarizes the key findings that inform the personality maturation principle.



Figure 11.2: Summary of the personality maturation principle.

While the maturity principle suggests a general trend of personality maturation as individuals age, there is one caveat. Researchers have found that the biological, social, and psychological changes that occur during the transition from childhood to adolescence can lead to temporary dips in some agreeableness, conscientiousness, and emotional stability (Tacket & Soto, 2015). For most, puberty is the period of life characterized by the most turbulent personality change. This idea of *personality disruption* will come as no surprise to any parents of teenagers reading this.

Normative Trends Aren't "One Size Fits All"

We've examined the major normative trends in personality development: personality becomes more stable as we age, and people tend to become more agreeable, emotionally stable, and conscientious. However, these trends reflect the *average* tendencies of how a population of people can be expected to change. They do not necessarily predict how any given individual will change—it's very possible that your own personality development doesn't seem to fit very well with these normative trends.

Perhaps the most significant finding of the last few decades of empirical research on personality change is that personality development is massively *heterogenous*, meaning that there are many individual differences in how personality changes. This heterogeneity is itself an interesting phenomenon that personality researchers are only beginning to study. Eventually, we may have a better sense of factors that predict differences in personality change. For now, it is simply important to recognize that, even though the empirical research has revealed robust trends in how populations tend change, normative trends are not "one size fits all". Personality changes differently for different people.

References

- Bleidorn, W. (2015). What accounts for personality maturation in early adulthood?. *Current Directions in Psychological Science*, 24(3), 245-252.
- Bleidorn, W., Schwaba, T., Zheng, A., Hopwood, C. J., Sosa, S. S., Roberts, B. W., & Briley, D. A. (2022). Personality stability and change: A meta-analysis of longitudinal studies. *Psychological bulletin*, 148(7-8), 588.
- Donnellan, M. B. (2023). Personality stability and change. In R. Biswas-Diener & E. Diener (Eds), *Noba textbook series: Psychology*. Champaign, IL: DEF publishers. Retrieved from <http://noba.to/sjvtxbwd>
- Shiner, R. L., Soto, C. J., & De Fruyt, F. (2021). Personality assessment of children and adolescents. *Annual Review of Developmental Psychology*, 3, 113-137.
- Soto, C. J., & John, O. P. (2014). Traits in transition: the structure of parent-reported personality traits from early childhood to early adulthood. *Journal of Personality*, 82(3), 182-199.
- Soto, C. J., & Tackett, J. L. (2015). Personality traits in childhood and adolescence: Structure, development, and outcomes. *Current Directions in Psychological Science*, 24(5), 358-362.
- Tackett, J. L., Slobodskaya, H. R., Mar, R. A., Deal, J., Halverson Jr, C. F., Baker, S. R., ... & Besevegis, E. (2012). The hierarchical structure of childhood personality in five countries: Continuity from early childhood to early adolescence. *Journal of Personality*, 80(4), 847-879.

12 Culture and Personality

You may not have realized this, but most of our exploration of personality thus far has been focused on a very small subset of humans. Most of the data is focused on people living in the United States, Canada, and the United Kingdom because that's where most psychological research has historically been conducted. Humans in these countries, however, are obviously only a very small percentage of the global population, and they are probably different in several regards from other humans who are currently living or have lived in the past. Specifically, these countries are geographically (and ideologically) *Western*, highly educated (in terms of formal schooling), heavily *Industrialized*, financially *Rich* (in terms of GDP), and *Democratic*. Psychologists commonly refer to this group of country characteristics with the tongue-in-cheek acronym "WEIRD".

This overrepresentation of people from WEIRD societies is a problem because most psychologists probably aren't just interested in the psychology of WEIRD people. We are interested in humans and human nature more broadly. To make more generalizable claims about human nature and to understand how human psychology and behavior works, we need study more diverse groups of people. That's where cross-cultural research comes in.

Cross-cultural research plays a pivotal role in enriching our understanding of human behavior and psychology by exploring the breadth of cultural variation. It allows researchers to discern the universality of psychological phenomena and discern the extent to which cultural factors shape them. In applied settings, cross-cultural research can inform the development of interventions and treatments that are culturally sensitive and suitable for diverse populations. At a societal level, learning about cross-cultural research can promote a more nuanced understanding and respect for diverse cultures. Additionally, it facilitates collaboration and exchange between researchers from different countries and cultural backgrounds, thereby enhancing the collective knowledge base of psychology on a global scale and fostering inclusivity and equity in psychological research.

In this chapter, we will examine some cross-cultural research exploring the interplay between culture and personality. We will examine what culture is, what it is made up of, and learn about a few fundamental dimensions of culture that will orient our focus. Then, we will explore how these fundamental cultural dimensions relate to cultural variation in personality trait levels and the structure of personality itself. Before we get going with our exploration of culture and personality, let's review some basic pitfalls that we want to avoid and review basics of statistical distributions that can help us avoid them.

Common Pitfalls to Avoid

When examining cultural differences, there are several pitfalls to avoid. **Essentialism** is a common pitfall to avoid when discussing cross-cultural research and cultural differences. This involves reducing a culture to a set of fixed, immutable traits or characteristics, ignoring the diversity and complexity within that culture. It overlooks the dynamic nature of cultures and individuals within them, leading to oversimplification and stereotyping. **Ethnocentrism** is another pitfall, where one evaluates other cultures solely based on the standards and values of their own culture, often leading to judgment or bias against those cultures. It's crucial to recognize and appreciate cultural differences without imposing one's own cultural norms as the universal standard.

We also want to avoid making a set of twin errors. The first is **minimizing differences between cultures that are not our own** which is problematic as it disregards the unique experiences, perspectives, and challenges faced by different cultural groups. Second, we want to avoid **exaggerating differences between cultures**, which can lead to misunderstandings and reinforce stereotypes, overlooking the commonalities and shared humanity that exist across cultures. Striking a balance between acknowledging cultural diversity and recognizing our commonalities is essential for meaningful cross-cultural understanding and research.

In my opinion, these pitfalls can be avoided by considering that cultural differences are usually referring only to the *average* of a distribution. Let's review the concepts of averages and data distributions that we learned about way back in the chapter on Foundational Frameworks.

Understanding how an average relates to an underlying distribution is crucial for putting cultural differences into context. When we talk about an average, we are essentially summarizing a set of data points into a single value. It's important to remember that this average is just one point within a broader distribution of data that captures the full variability or spread of individual scores within a group or culture.

Research across group differences consistently shows that there is more variation *within* groups than *between* groups. If we were to measure extraversion of people living in the USA and China, for example, we might find an average difference in extraversion between the countries; but within each country, there would still be a larger amount of variation among the people than the average difference between the countries. In other words, while there may be variability in averages between cultures, these differences are often small in magnitude when viewed in the context of the overall variability within each culture.

From the perspective of effect size measures like Cohen's d , which quantify the magnitude of differences between groups, average differences between cultures tend to be relatively small ($d < .2$; for context, the difference in average height of men vs. women is around seven times bigger). This means that there is usually substantial overlap between groups or cultures in terms of the distribution of scores on a given trait or characteristic. Therefore, it's essential not to overstate the significance of average differences between cultures and to recognize the diversity and variability within each culture.

By acknowledging the overlap between cultural groups, we can avoid the trap of oversimplifying or stereotyping entire cultures based on average differences. Instead, we can appreciate the richness and complexity of human diversity, recognizing that individuals within any culture still exhibit a wide range of behaviors, attitudes, and traits. This perspective fosters a more nuanced understanding of cultural differences and promotes respectful and inclusive cross-cultural dialogue and research.

Now that we are equipped to avoid these pitfalls, we can begin to traverse the landscape of cultural variation in personality.

Elements of Culture

What is culture? Just as it is difficult to pin down just one definition of personality, it is also challenging to define culture succinctly. For our purposes, we will avoid providing a rigid definition of culture and instead focus on exploring some of the key elements that constitute it. We'll consider five elements of culture: symbols, beliefs, values, norms, and practices.

Symbols refer to the representations of language or thoughts that have shared meaning within a group. Symbols can be just about anything that has an agreed upon meaning. Some examples include: the alphabet letters that make up the words that allow me to communicate these concepts to you right now; mathematical notation that represents algorithmic procedures and abstract numbers; and emojis.

Beliefs refer to what people in a culture hold to be true, reflecting their understanding of the world. Examples of beliefs include religious beliefs, such as the belief in an afterlife or the existence of a higher power; scientific beliefs, such as the belief in the theory of evolution or the laws of physics; and cultural beliefs, such as the belief in the importance of hospitality or the power of superstitions.

Values are deeply ingrained principles that guide individuals' actions and judgments, reflecting what is considered important or desirable within a cultural context. Examples of values include cultural values such as respect for elders or the pursuit of individual happiness; moral values such as honesty, integrity, or compassion; and societal values such as democracy, freedom, or justice.

Norms dictate what behaviors and interactions are expected or typical within a society. They can serve as social guidelines that promote cohesion and cooperation. Examples of norms include social norms such as queuing in line or waiting for one's turn; gender norms such as expectations regarding clothes, appearance, or behavior based on gender; and legal norms such as laws against theft, violence, or discrimination.

Practices encompass the tangible expressions of culture, including rituals, ceremonies, and customs that are enacted by individuals within a cultural group. Examples of practices include religious rituals such as prayer, meditation, or worship services; cultural ceremonies such as

weddings, funerals, or coming-of-age celebrations; and everyday customs such as greetings, gestures, or dietary habits.

Together, these elements interact in different combinations to make up the culture of a group. These elements are not totally independent; they often influence each other. The *practice* of praying, for instance, likely depends on the *belief* that an entity can answer prayers.

While there may be an underlying logic to many combinations of cultural elements that arise in a group, there can also be a randomness or arbitrariness to it. Some cultural elements, such as norms or practices, may have originated by chance or historical accidents; but they are perpetuated simply because people continue to adhere to them. Culture is dynamic and evolving — shaped by a complex interplay of historical, social, and psychological factors. There is a whole field of cultural evolution that specifically studies the dynamic change of cultural elements over time.

Levels of Culture and Cultural Identities

It is also important to recognize that culture can be defined at many levels. A nation is a convenient unit for cross-cultural research, but people also belong to ethnic groups, political groups, social classes, cities, neighborhoods, religions, languages, occupations, hobbies, music scenes, academic disciplines, and many other groups. Each of these can include its own symbols, beliefs, values, norms, and practices.

This means that a person is not simply from one culture. Most people hold many cultural identities at once, and these identities can become more or less important across situations. For example, the cultural expectations that guide behavior in a family setting may not be the same expectations that guide behavior in a classroom, workplace, religious community, neighborhood, or online group. This is one reason we should expect cultural influences on personality to be plural, probabilistic, and heterogeneous rather than simple and deterministic.

Dimensions of Culture

Just like there are countless ways individual people might differ in their tendencies of thoughts, feelings, and behaviors (i.e., personality traits), there are also countless ways that groups can differ in their unique combinations of symbols, practices, beliefs, norms, and values (i.e., culture). In previous chapters, we learned that psychologists made the study of personality traits more tractable by using factor analysis to identify broad dimensions of personality; it turns out that a similar approach was applied to identify broad dimensions of culture.

In the 1960s and 1970s Dr. Geert Hofstede, a psychologist studying cultural differences, used interviews and surveys to index variation in practices, beliefs, norms, and other elements of culture across multiple countries. Through his research and several follow-up studies, six fundamental cultural dimensions have been identified (Hofstede & Bond, 1984). The first

dimension, *Power Distance*, reflects the acceptance of unequal power distribution within a society. *Individualism vs. Collectivism*, the second dimension, measures the prioritization of personal versus group interests. The third dimension, *Masculinity vs. Femininity*, explores the degree of emphasis on traditional masculine values like competitiveness and assertiveness versus feminine values such as cooperation and nurturance. *Uncertainty Avoidance*, the fourth dimension, examines the level of discomfort with ambiguity and risk within a culture. The fifth dimension, introduced later, is *Long-Term Orientation vs. Short-Term Orientation*, which assesses the focus on long-term planning and perseverance versus short-term gratification. Finally, *Indulgence vs. Restraint*, indexes the extent to which societies allow for gratification of human desires (e.g., sexual, material). See the table below for an overview of what the extremes of each dimension describe.

Table 12.1: Description of Hofstede’s six dimensions of culture

Dimension	Description of High End	Description of Low End
Power Distance	Acceptance of hierarchical structures and centralized decision-making. Less powerful members of society defer to authority figures and expect clear delineation between social classes or ranks.	Preference for more egalitarian relationships. Individuals challenge authority and expect a participatory decision-making process.
Individualism vs. Collectivism	Emphasis on personal freedom, autonomy, and achievement. People prioritize their own goals and assert their individuality.	Emphasis on group harmony and cooperation. Individuals prioritize the needs and goals of the group over individual desires.
Masculinity vs. Femininity	Emphasis on traditionally masculine values such as competitiveness, assertiveness, and achievement. Success is often defined in terms of material wealth and status.	Emphasis on nurturing relationships, quality of life, and work-life balance. Success is often defined in terms of social welfare and personal fulfillment.
Uncertainty Avoidance	Preference for rules, structure, and conformity to minimize uncertainty. People avoid ambiguous situations and seek security through adherence to traditions and established norms.	Tolerance for ambiguity and a willingness to take risks. People are more open to change, innovation, and experimentation.

Dimension	Description of High End	Description of Low End
Long-Term Orientation vs. Short-Term Orientation	Emphasis on long-term planning, perseverance, and thriftiness. People prioritize long-term goals and are willing to sacrifice short-term gains for future success.	Emphasis on short-term gratification, tradition, and respect for social norms. People focus on enjoying the present moment and may prioritize short-term gains over long-term planning.
Indulgence vs. Restraint	Emphasis on enjoying life, freedom of expression, and indulging in personal desires. People prioritize leisure, pleasure, and self-expression.	Emphasis on self-discipline, modesty, and adherence to social norms. People restrain their desires and adhere to strict social codes of conduct.

Importantly, much like the Big Five and other factor analytic models of personality do not capture every nuance of personality variation, Hofstede's dimensions also probably do not capture many nuances of cultural differences. However, the dimensions provide a tractable framework for understanding and comparing cultural differences across societies. We will refer to these dimensions as we further explore the interplay between culture and personality.

Interplay Between Culture and Personality

How is culture related to personality? As you might guess, these constructs are overlapping, and mutually reinforcing. As we will see, it is likely that personality shapes culture (personality $\rightarrow$ culture), and that culture shapes personality (culture $\rightarrow$ personality). This is because the systems of interconnected beliefs, norms, values, and other elements of culture are ultimately maintained in the minds of individuals: individuals are culturally influenced influencers of culture. But also, culture can shape personality via at least two pathways: enculturation and acculturation.

Enculturation is the process through which individuals learn and acquire the cultural norms, values, beliefs, and behaviors of the society they are born into or grow up in. This process begins at a very young age and continues throughout one's life, shaping identity and sense of belonging. Enculturation occurs primarily through socialization, which involves interactions with family members, peers, educators, media, and other social institutions. For example, a child raised in a culture that values collectivism may be taught to prioritize group harmony and cooperation over individual achievement, while a culture that emphasizes independence may encourage children to assert their autonomy and self-reliance.

Acculturation, on the other hand, refers to the process of adopting and incorporating elements of a different culture into one's own cultural identity. This often occurs when individuals

or groups encounter a different culture due to migration, colonization, or globalization. Acculturation involves both cultural exchange and adaptation, as individuals may selectively adopt certain aspects of the new culture while retaining elements of their original cultural identity. For instance, immigrants to a new country may adopt certain customs, language, and practices of the host culture while still maintaining their native language, cuisine, and traditions. Over time, acculturation can lead to the development of a new blends of cultural practices and identities that reflect the diverse experiences and backgrounds of individuals or groups.

There are some important nuances to keep in mind when considering the power of these cultural processes to shape personality. Firstly, they are implicit, meaning that individuals often internalize cultural norms and values without conscious awareness. Secondly, they are plural, indicating that there are multiple cultural influences at play, including family, peers, media, and societal institutions. Thirdly, they are probabilistic, suggesting that individuals may not adopt all aspects of their culture equally or in the same way, and outcomes can vary based on various factors such as individual differences and environmental circumstances. Finally, these processes are heterogeneous, meaning that they manifest differently across diverse cultural contexts and can result in a wide range of outcomes in terms of personality, beliefs, and behaviors. With these caveats in mind, let's examine some of the empirical research describing how cultural influences may shape personality.

Culture → Personality

One way culture can shape personality is by repeatedly exposing people to particular situations, expectations, rewards, and punishments. A culture that strongly values group harmony, for example, may provide frequent opportunities to practice deference, cooperation, and concern for others. A culture that strongly values self-expression may provide more opportunities to practice assertiveness, novelty seeking, and public self-promotion. These repeated situations do not create personality out of nothing, but they may help calibrate the average ways that people think, feel, and behave.

Cross-cultural research has found average differences in Big Five trait levels across nations and cultural groups (Schmitt et al., 2007). These findings are interesting, but they should be interpreted with the cautions from earlier in the chapter. Average differences in trait levels do not imply that every person in one culture differs from every person in another culture. They also do not tell us whether culture caused the difference, personality caused the cultural pattern, both influenced each other, or some third variable influenced both.

Culture and Personality Development

Culture may also shape how personality changes across the lifespan. Recall from the development chapter that people tend, on average, to become more agreeable, conscientious, and emotionally stable with age. Cross-cultural research suggests that these maturation trends are

fairly widespread, but the timing and strength of these changes differ across cultures. One explanation is that cultures differ in when people enter adult social roles, such as stable work, committed partnerships, parenting, and other responsibilities. In one large study across 62 nations, cultures with earlier onset of adult-role responsibilities tended to show earlier personality maturation (Bleidorn et al., 2013).

Cultural expectations may also matter later in life. For example, Kitayama and colleagues (2020) argued that some Western cultural contexts emphasize remaining youthful, active, positive, and independent even in late adulthood. That cultural expectation may create a mismatch for older adults as physical and social conditions change. In contrast, some East Asian cultural contexts place more emphasis on age-graded roles and adapting to changing social positions across the lifespan.

These patterns suggest that culture may influence which personality patterns are reinforced, expected, or made difficult at different ages.

Cultural Experiences and Acculturation

Acculturation provides another window into culture shaping personality. When people move into a new cultural context, they may repeatedly encounter different norms, values, practices, and social expectations. Over time, these experiences can shift how people typically behave or how they understand themselves. Research on studying abroad, for example, suggests that international experiences can be associated with changes in some Big Five traits, although people who choose to study abroad may already differ from people who do not (Niehoff et al., 2017).

Research on immigrants provides a similar example. In one study of Japanese, Japanese American, and European American participants, greater involvement in U.S. culture was associated with personality profiles that fit U.S. cultural patterns more closely and Japanese cultural patterns less closely (Güngör et al., 2013). This is consistent with the idea that personality can acculturate. But it still requires caution: cultural exposure, self-selection, age, language, social roles, and measurement issues can all be tangled together.

Problems with Comparing Personality Across Cultures

Comparing personality across cultures is difficult because the measure itself may not work the same way across groups. A translated item may not carry the same meaning. A trait word may depend on a cultural context that is not shared by all participants. People in different cultures may compare themselves to different reference groups when answering questions (Heine et al., 2002). Response styles can also differ; for example, some groups may be more likely to agree with statements in general, avoid extreme responses, or answer in socially desirable ways.

These problems matter because they can make a real cultural difference look larger or smaller than it is, or they can make a measurement artifact look like a cultural difference. This does not mean we should give up on cross-cultural personality research. It means we should treat reliability and validity as central scientific problems rather than boring technical details.

Personality → Culture

Culture can shape personality, but the arrow also points in the other direction. Cultures are produced and maintained by groups of individuals. If many people in a group tend to be open to novelty, extraverted, threat-sensitive, cooperative, suspicious, risk tolerant, or highly status motivated, those aggregate tendencies may influence the symbols, beliefs, values, norms, and practices that become common in that group.

For example, region-level personality has been linked to cultural and political outcomes. Obschonka and colleagues (2018) found that regional levels of neurotic traits were associated with support for Brexit in the United Kingdom and Donald Trump in the United States. This kind of finding does not mean that a single trait simply explains voting behavior. It does show that aggregate personality can be connected to large-scale cultural outcomes, especially when social conditions make certain concerns, motives, or emotions politically relevant.

The difficulty is that these patterns are usually circular. Personality may help shape culture, culture may help shape personality, and both may be influenced by other features of the environment. A cautious population might develop cautious norms, and cautious norms might then reinforce cautious behavior. Over time, it becomes hard to separate the chicken from the egg.

Ecological Factors

One way to make progress on the chicken-and-egg problem is to examine ecological factors. Ecological factors are aspects of the physical environment that can influence development and behavior but are less likely to be directly produced by personality or culture. Examples include latitude, temperature, topography, pathogen prevalence, soil quality, resource availability, and subsistence style. These variables are not magic causal explanations, but they can help us think more clearly about why cultures and personality patterns might differ across places.

Consider subsistence style. Some forms of farming require more coordination than others. Wetland rice farming, for example, often requires coordinated irrigation, shared labor, and monitoring of water use. Talhelm and English (2020) found that historically rice-farming regions tended to have tighter social norms than wheat-farming regions, both within China and across nations. This is not saying that rice causes one personality type. It is a hypothesis about how repeated ecological demands can shape cultural norms, which may then shape behavior and personality expression.

Pathogen prevalence provides another example. In environments where infectious disease risk was historically high, norms that limited contact with unfamiliar outsiders or emphasized group boundaries may have reduced exposure to disease. Consistent with this logic, pathogen prevalence has been linked to cultural variation in individualism-collectivism and to average levels of traits such as extraversion and openness (Fincher et al., 2008; Schaller & Murray, 2008). Again, this is not a simple deterministic story. Pathogens, economy, institutions, climate, migration, and measurement issues can all be intertwined.

Topography is another possible ecological influence. Mountainous or frontier environments may be more difficult to settle and may reward or attract people who are more independent, self-reliant, or open to risk and novelty. In a large study of the United States, Götz and colleagues (2020) found that mountainousness was associated with regional differences in Big Five traits. These effects were not enormous, and they do not mean that every person in a mountain region has the same personality. But they illustrate how physical environments can become part of the broader explanation for cultural and personality variation.

Cultural Variation in Personality Structure

So far, we have mostly focused on trait levels: whether one group is higher or lower, on average, on some personality trait. But we also need to think about personality structure. Personality structure refers to the pattern of covariation among traits in a population. The Big Five is not merely a list of five interesting traits. It is a claim that many narrower traits tend to cluster into five broad dimensions.

Many trait theorists have argued that the Big Five or other “Big Few” models reflect basic biological tendencies shared by humans. If that view is correct in a strong form, then the same broad structure should appear across cultures. We might expect different cultures to vary in average trait levels, but the underlying trait structure should remain basically the same.

The evidence is more complicated. The Big Five structure tends to replicate fairly well in many Western, educated, industrialized, rich, and democratic populations. But the structure is often less clear outside these contexts, especially in some small-scale or non-industrialized societies (Gurven et al., 2013; McCrae & Terracciano, 2005; Schmitt et al., 2007). This raises an important question: Is personality structure universal, or does it partly depend on socioecological context?

Another way to see this issue is through indigenous personality research. If researchers begin with personality concepts that are especially important in a particular language or cultural context, they may find dimensions that imported Big Five measures miss. For example, research on Chinese personality assessment has identified an Interpersonal Relatedness dimension, including concepts such as *ren qing*, or reciprocal relationship orientation, that are not fully captured by standard Big Five measures (Cheung et al., 2001). This does not mean people in one culture have psychological traits that people in another culture completely lack. It means

the socially important ways of describing, organizing, and measuring personality may differ across cultures.

The Tsimane Case Study

One important case study comes from research with the Tsimane, an Indigenous group living in the Bolivian Amazon. The Tsimane are often described as forager-horticulturalists who live in relatively small communities, rely heavily on subsistence activities, and have had less exposure to formal schooling and industrialized institutions than participants in typical personality studies.

Gurven and colleagues (2013) translated a Big Five questionnaire into the Tsimane language and interviewed more than 600 adults. They also collected spouse reports from another sample. Across these analyses, the Big Five structure did not clearly emerge. Instead, the data were more consistent with two broad dimensions: prosociality and industriousness. These dimensions mixed together narrower traits that would usually be split across different Big Five dimensions in WEIRD samples.

There are two ways to react to this finding. One reaction is to assume the result is just a measurement problem: maybe the translation was imperfect, the questions were unfamiliar, or the response format did not fit the context. Those are serious possibilities, and cross-cultural researchers should take them seriously. But Gurven and colleagues tested several measurement explanations and still did not find strong support for the Big Five. So the more interesting possibility is that the structure of personality itself may differ across socioecological contexts.

The Niche Diversity Hypothesis

The niche diversity hypothesis was developed to explain why personality structure might differ across societies. A niche is a smaller position or role within a larger society that has its own opportunities, constraints, and payoff structures. Occupations, social groups, coalitions, families, and organizations can all be thought of as niches. Different niches may reward different combinations of traits.

In a society with many diverse niches, there may be many different ways to be successful. One niche might reward patience, diligence, and emotional stability. Another might reward sociability and risk tolerance. Another might reward technical skill and low concern with public attention. If people calibrate their behavior to the niches they occupy, then different combinations of traits can become useful in different parts of society. At the population level, this should reduce overlap among traits and produce more distinct personality dimensions (Lukaszewski et al., 2017; Smaldino et al., 2019).

In a society with fewer distinct niches, the same combinations of traits may be useful across more domains of life. If success in domestic, cooperative, productive, and public activities depends on overlapping trait combinations, then traits that appear separate in one society may covary more strongly in another. This could produce fewer broad dimensions, such as the prosociality and industriousness pattern observed among the Tsimane.

Smaldino and colleagues (2019) tested the logic of this idea with agent-based computer simulations. The simulations showed that if people occupy niches, niches reward particular trait profiles, and people adjust toward the traits rewarded by their niches, then populations with more niches develop lower correlations among traits and more personality dimensions. A simulation does not prove that the hypothesis is true in real humans, but it shows that the hypothesis is logically coherent.

Durkee and colleagues (2022) tested the niche diversity hypothesis using Big Five data from 685,089 people across 115 nations. Niche diversity was estimated using indicators such as human development, urbanization, and economic complexity. Countries with higher niche diversity tended to show lower covariance among personality traits and greater personality dimensionality. In simpler terms, countries with more diverse niches tended to have more distinct personality dimensions. The study did not find consistent evidence that niche diversity was related to wider trait distributions.

These findings suggest that niche diversity can partly explain cultural variation in personality structure. But it is not without problems. The evidence is correlational, the indicators of niche diversity are imperfect, the personality data came from people with internet access, and cross-cultural personality measurement remains difficult. Still, this framework helps move the field beyond simply asking whether the Big Five is universal. A better question is: What features of human minds, cultures, and socioecologies explain both the similarities and differences in personality structure across populations?

Closing Thoughts About Culture and Personality

Culture and personality are not separate forces acting independently on behavior. Culture is maintained in the minds and actions of individuals, and individuals develop within cultural and ecological contexts. This means that culture can shape personality, personality can shape culture, and ecological factors can influence both.

The goal of this chapter is not to throw out trait models like the Big Five or HEXACO. Those models are useful tools for summarizing personality variation. The goal is to recognize their limits. A complete science of personality has to explain not only why some patterns appear across many human groups, but also why those patterns change across cultures, ecological contexts, historical periods, and social niches.

References

- Bleidorn, W., Klimstra, T. A., Denissen, J. J. A., Rentfrow, P. J., Potter, J., & Gosling, S. D. (2013). Personality maturation around the world: A cross-cultural examination of social-investment theory. *Psychological Science*, 24(12), 2530-2540. <https://doi.org/10.1177/0956797613498396>
- Cheung, F. M., Leung, K., Zhang, J. X., Sun, H. F., Gan, Y. Q., Song, W. Z., & Xie, D. (2001). Indigenous Chinese personality constructs: Is the five-factor model complete? *Journal of Cross-Cultural Psychology*, 32(4), 407-433. <https://doi.org/10.1177/0022022101032004003>
- Durkee, P. K., Lukaszewski, A. W., von Rueden, C. R., Gurven, M. D., Buss, D. M., & Tucker-Drob, E. M. (2022). Niche diversity predicts personality structure across 115 nations. *Psychological Science*, 33(2), 285-298. <https://doi.org/10.1177/09567976211031571>
- Fincher, C. L., Thornhill, R., Murray, D. R., & Schaller, M. (2008). Pathogen prevalence predicts human cross-cultural variability in individualism/collectivism. *Proceedings of the Royal Society B: Biological Sciences*, 275(1640), 1279-1285. <https://doi.org/10.1098/rspb.2008.0094>
- Götz, F. M., Stieger, S., Gosling, S. D., Potter, J., & Rentfrow, P. J. (2020). Physical topography is associated with human personality. *Nature Human Behaviour*, 4, 1135-1144. <https://doi.org/10.1038/s41562-020-0930-x>
- Güngör, D., Bornstein, M. H., De Leersnyder, J., Cote, L., Ceulemans, E., & Mesquita, B. (2013). Acculturation of personality: A three-culture study of Japanese, Japanese Americans, and European Americans. *Journal of Cross-Cultural Psychology*, 44(5), 701-718. <https://doi.org/10.1177/0022022112470749>
- Gurven, M., von Rueden, C., Massenkoff, M., Kaplan, H., & Lero Vie, M. (2013). How universal is the Big Five? Testing the five-factor model of personality variation among forager-farmers in the Bolivian Amazon. *Journal of Personality and Social Psychology*, 104(2), 354-370. <https://doi.org/10.1037/a0030841>
- Heine, S. J., Lehman, D. R., Peng, K., & Greenholtz, J. (2002). What's wrong with cross-cultural comparisons of subjective Likert scales? The reference-group effect. *Journal of Personality and Social Psychology*, 82(6), 903-918. <https://doi.org/10.1037/0022-3514.82.6.903>
- Hofstede, G., & Bond, M. H. (1984). Hofstede's culture dimensions: An independent validation using Rokeach's value survey. *Journal of Cross-Cultural Psychology*, 15(4), 417-433.
- Kitayama, S., Berg, M. K., & Chopik, W. J. (2020). Culture and well-being in late adulthood: Theory and evidence. *American Psychologist*, 75(4), 567-576. <https://doi.org/10.1037/amp0000614>
- Lukaszewski, A. W., Gurven, M., von Rueden, C. R., & Schmitt, D. P. (2017). What explains personality covariation? A test of the socioecological complexity hypothesis. *Social Psychological and Personality Science*, 8(8), 943-952. <https://doi.org/10.1177/1948550617697175>

- McCrae, R. R., & Terracciano, A. (2005). Universal features of personality traits from the observer's perspective: Data from 50 cultures. *Journal of Personality and Social Psychology*, 88(3), 547-561. <https://doi.org/10.1037/0022-3514.88.3.547>
- Niehoff, E., Petersdotter, L., & Freund, A. P. (2017). International sojourn experience and personality development: Selection and socialization effects of studying abroad and the Big Five. *Personality and Individual Differences*, 112, 55-61. <https://doi.org/10.1016/j.paid.2017.02.043>
- Obschonka, M., Stuetzer, M., Rentfrow, P. J., Lee, N., Potter, J., Gosling, S. D., & Schmitt-Rodermund, E. (2018). Fear, populism, and the geopolitical landscape: The sleeper effect of neurotic personality traits on regional voting behavior in the 2016 Brexit and Trump votes. *Social Psychological and Personality Science*, 9(3), 285-298. <https://doi.org/10.1177/1948550618755874>
- Schaller, M., & Murray, D. R. (2008). Pathogens, personality, and culture: Disease prevalence predicts worldwide variability in sociosexuality, extraversion, and openness to experience. *Journal of Personality and Social Psychology*, 95(1), 212-221. <https://doi.org/10.1037/0022-3514.95.1.212>
- Schmitt, D. P., Allik, J., McCrae, R. R., Benet-Martinez, V., Alcalay, L., Ault, L., Austers, I., Bennett, K. L., Bianchi, G., Boholst, F., Borg Cunen, M. A., Braeckman, J., Brainerd, E. G., Caral, L. G. A., Caron, G., Casullo, M. M., Cunningham, M., Daibo, I., De Backer, C., ... Zupancic, A. (2007). The geographic distribution of Big Five personality traits: Patterns and profiles of human self-description across 56 nations. *Journal of Cross-Cultural Psychology*, 38(2), 173-212. <https://doi.org/10.1177/0022022106297299>
- Smaldino, P. E., Lukaszewski, A., von Rueden, C., & Gurven, M. (2019). Niche diversity can explain cross-cultural differences in personality structure. *Nature Human Behaviour*, 3, 1276-1283. <https://doi.org/10.1038/s41562-019-0730-3>
- Talhelm, T., & English, A. S. (2020). Historically rice-farming societies have tighter social norms in China and worldwide. *Proceedings of the National Academy of Sciences*, 117(33), 19816-19824. <https://doi.org/10.1073/pnas.1909909117>

13 Personality Change

Earlier in the course, we examined personality development across the life course. We learned that personality traits are not perfectly fixed. People tend to become more stable in their relative standings as they get older, and there are average developmental trends in traits such as agreeableness, conscientiousness, and emotional stability. But that still leaves a more personal and practical question: can people change their personality in the specific ways they want to change?

This is a tricky question because the answer depends on what we mean by change. If by change we mean becoming a completely different person overnight, then the answer is almost certainly no. Personality traits are relatively stable summaries of typical patterns of thoughts, feelings, and behaviors. They do not usually flip all at once. But if by change we mean shifting the average pattern of thoughts, feelings, and behaviors over weeks, months, or years, then the answer appears to be yes. The more difficult question is how this happens, how much change is possible, and what conditions make change more likely.

In this chapter, we will examine volitional personality change, which refers to intentional efforts to change one's own personality traits. We will also examine evidence from clinical interventions, because psychotherapy and other structured interventions may change personality even when personality change is not the explicit goal. Along the way, we will keep returning to a basic idea: traits are broad descriptions, so changing a trait probably requires repeatedly changing the smaller state-level thoughts, feelings, and behaviors that make up the trait.

Many People Want to Change Their Personality

Before asking whether people can change their traits, it is worth noting that many people want to. If you ask people whether there is some aspect of their personality they would like to change, most people can think of something. This should not be surprising. Personality traits are tied to many outcomes people care about, including relationships, work, health, and emotional well-being.

The most common desired changes also make intuitive sense. Across studies, people often report wanting to become more emotionally stable, more conscientious, more extraverted, and more agreeable (Baranski et al., 2021; Hudson & Fraley, 2015). In everyday terms, many people would like to feel less anxious or emotionally reactive, become more disciplined and organized, feel more socially confident, or be kinder and more patient with other people. These desired

changes overlap with the maturity principle from the development chapter, which described the average tendency for people to become more agreeable, conscientious, and emotionally stable with age.

However, wanting to change is not the same thing as changing. Lots of people want to exercise more, sleep better, save money, or study earlier for exams. The wanting is real, but it does not automatically produce the behavior. Personality change has the same problem. The question is not only whether people prefer a different trait profile. The question is whether they can repeatedly think, feel, and behave in ways that are different enough, and consistent enough, to eventually shift their average tendencies.

From Personality States to Personality Traits

A useful way to think about personality change is to distinguish between personality states and personality traits. A personality state is a temporary pattern of thinking, feeling, or behaving in a particular moment. For example, a person might behave assertively during a group task, conscientiously while studying, or (dis)agreeably during a conflict. A personality trait is a broader summary of how often and how strongly a person tends to show those states across many situations and over time.

This distinction makes personality change seem less mysterious. To become more extraverted, a person does not need to magically acquire an extraversion substance inside the brain. They would need to spend more of their life in extraverted states: initiating conversations, expressing enthusiasm, approaching social opportunities, or practicing social confidence. To become more conscientious, a person would need to spend more of their life in conscientious states: planning, following through, organizing tasks, and resisting impulses that interfere with goals.

The basic premise of many personality change interventions is that lasting trait change can be produced by modifying state-level thoughts, feelings, and behaviors over extended periods of time (Hudson, 2021; Hudson et al., 2021). This does not mean that every temporary state becomes a trait. Acting outgoing one time will not make someone an extravert. Studying carefully for one exam will not make someone highly conscientious. The idea is that repeated states, practiced over time, may gradually become easier, more automatic, and more incorporated into the person's identity.

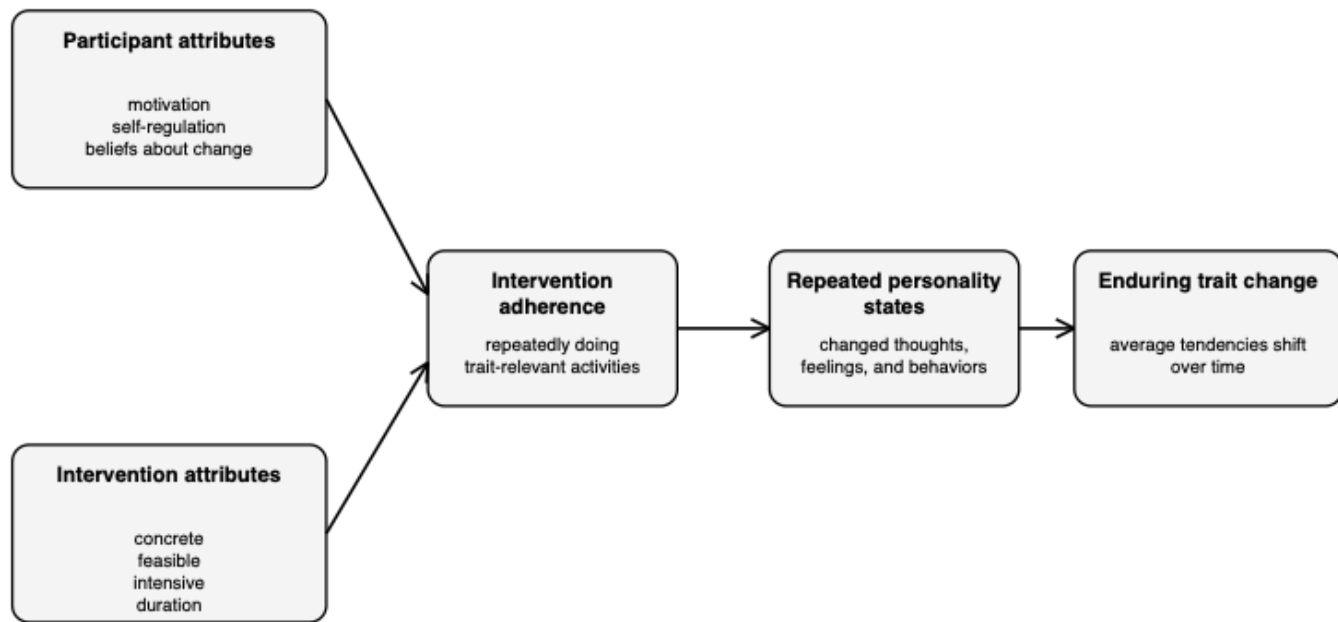


Figure 13.1: A simplified framework for thinking about personality change interventions.

Figure 13.1 summarizes one way to think about this process. Interventions do not directly reach into the mind and rewrite a personality trait. Instead, interventions are supposed to change what people repeatedly do. The most immediate bottleneck is intervention adherence, which means whether the person actually follows through with the activities, challenges, or behavioral practices that the intervention requires. If adherence is low, then the intervention has little opportunity to change the person's recurring states. If adherence is high, then those recurring states may eventually contribute to identity changes, biological changes, and trait change.

This is why personality change is not just a matter of insight. Insight may help, and motivation may help, but the central mechanism in this framework is repeated change in thoughts, feelings, and behaviors. If the behavioral pattern never changes, then there is not much for the trait summary to summarize differently.

Participant Attributes That May Affect Change

People differ in how likely they are to benefit from a personality change intervention. One obvious factor is motivation. A person who strongly wants to change may be more willing to practice new behaviors, tolerate discomfort, and continue when progress is slow. Motivation may be especially important because many personality change goals require doing things that

initially feel unnatural. A shy person trying to become more extraverted may need to start conversations that feel awkward. A disorganized person trying to become more conscientious may need to build routines that feel tedious.

Self-regulation is another important factor. Self-regulation refers to the ability to monitor one's behavior, inhibit competing impulses, maintain attention to goals, and adjust when things are not working. This matters because personality change is not just a single decision. It requires many small decisions across time. A person might decide in the morning to be more patient, more organized, or more outgoing, but the relevant test happens later when they are tired, stressed, busy, or distracted.

A third factor involves beliefs about change. Some people hold a relatively fixed mindset about personality, where they believe personality is basically set and cannot change much. Other people hold a growth mindset, where they believe personality is fluid and can change. It is easy to imagine that a growth mindset should make personality change more likely. If people believe change is possible, perhaps they will try harder and persist longer.

The evidence is less supportive of this idea than you might expect, however. In studies of emerging adults, beliefs about whether personality can change did not consistently predict actual trait change (Hudson et al., 2021). This does not mean beliefs are irrelevant in every possible context. It does mean that merely believing personality can change is not enough. A person can believe change is possible and still fail to follow through. Likewise, a person can be skeptical about personality change and still change if their daily behavior is repeatedly altered by an intervention, social role, or environment.

This finding is useful because it pushes us away from overly simple slogans. Growth mindsets may be encouraging, but personality change probably depends more on what people repeatedly do than on whether they endorse a general belief that change is possible.

Intervention Attributes That May Affect Change

The design of the intervention also matters. A vague goal such as 'be more conscientious' or 'be less anxious' may express a real desire, but it does not give the person much guidance. What should they do at 9:00 AM tomorrow? What should they do when a friend texts them, when they sit down to study, or when they feel irritated? Personality traits are broad summaries, but interventions need to operate at the level of concrete states and behaviors.

One important feature is *concreteness*. A concrete intervention translates a broad trait goal into specific, detailed actions. For example, instead of simply trying to be more organized, a person might decide to write a prioritized task list before bed each night, set a reminder for one recurring obligation, or spend ten minutes clearing a work space before starting homework. The point is not that these particular examples work for everyone. The point is that a concrete plan specifies what the person is supposed to do.

A second feature is *feasibility*. A goal can be concrete but unrealistic. If a person who currently avoids nearly all social interaction sets the goal of attending three large social events every week, the plan may fail because it is too demanding. Feasible interventions are challenging enough to create change but realistic enough that the person can actually follow through. This matters because failed plans can reduce motivation and make the person less likely to continue.

A third feature is *intensiveness*. An intervention that happens once is unlikely to produce lasting change unless it triggers repeated practice afterward. Personality traits summarize average patterns, so the intervention needs enough repetition to affect those averages. This is why many interventions involve repeated challenges, prompts, reminders, or exercises over weeks or months.

A fourth feature is *duration*. Even an intensive intervention may not last long enough for new behaviors to become habitual or identity-relevant. Duration refers to how long the intervention continues. Longer interventions are not automatically better, because people may stop adhering if the intervention becomes exhausting or inconvenient. But in general, changing a trait-like pattern should require sustained practice over time rather than a single burst of effort.

Evidence from Psychotherapy and Clinical Intervention

One source of evidence for personality change comes from psychotherapy and clinical interventions. These interventions are often designed to reduce distress, treat symptoms, or improve functioning, not necessarily to change personality traits. Even so, they may alter the recurring thoughts, feelings, and behaviors that personality questionnaires measure.

A meta-analysis of intervention studies found that personality traits can change during psychological and clinical interventions, with especially notable changes in emotional stability (Roberts et al., 2017). This is important because emotional stability is closely related to neuroticism, the tendency to experience frequent or intense negative emotions. If therapy reduces patterns of anxiety, worry, emotional reactivity, avoidance, or depressive thinking, then it is not surprising that people may also appear to become more emotionally stable on personality measures.

Psychotherapy is therefore a strong reminder that personality is not sealed off from the rest of psychology. The same processes that change symptoms, coping strategies, relationship patterns, or daily routines may also change trait-relevant patterns. Still, we should be cautious. When a person's depression or anxiety improves, their responses to personality items may change partly because they are in a different emotional state. Distinguishing true trait change from symptom change, response bias, or temporary state change is difficult. The evidence suggests that intervention-related trait change is possible, but it does not mean that every therapy client changes personality in the same way or to the same degree.

Evidence from Volitional Personality Change Studies

Research on volitional personality change asks a slightly different question: when people intentionally try to change their traits, do they move in the direction they want to move? Early studies suggest that they often do, at least to some degree. For example, people who reported wanting to increase a trait such as extraversion tended to show increases in that trait over time (Hudson & Fraley, 2015). This pattern is encouraging, but it also leaves open an important question: what are people doing that produces the change?

One answer comes from research on behavioral change challenges. In these studies, participants selected weekly challenges designed to pull their thoughts, feelings, and behaviors in line with their desired traits. For example, a challenge aimed at extraversion might involve initiating a conversation; a challenge aimed at conscientiousness might involve completing a task before a deadline. The key finding was that merely accepting challenges was not enough. Completing the challenges predicted trait change over time (Hudson et al., 2019).

This finding fits neatly with the state-to-trait logic described earlier. A challenge only matters if it changes what the person actually does. Accepting a challenge is a declaration of intent. Completing a challenge is a new behavior. Across enough repetitions, those new behaviors may accumulate into a different average pattern.

The evidence is not uniformly simple. Some interventions work better for some traits than others, and some attempts may produce weak or even unexpected results. Openness, for example, has sometimes been more difficult to change in a straightforward way than traits like extraversion, conscientiousness, agreeableness, or emotional stability. This should not be surprising. A personality trait is not a single behavior. It is a broad, heterogeneous construct. Some parts of a trait may be easier to practice than others, and some interventions may target only a narrow slice of the broader trait.

More recent intervention studies have used digital tools to support personality change. In one randomized controlled trial, adults used a smartphone-based coaching intervention designed to support desired trait changes over several months (Stieger et al., 2021). Participants receiving the intervention showed greater changes than a waitlist control group, and some changes were also visible to people who knew the participants. This provides stronger evidence than simple correlational studies, because random assignment helps rule out some alternative explanations. But even randomized interventions do not answer every question. We still need to know how long changes last, which components are most important, and whether changes produce improvements in life outcomes.

How Much Change Is Possible?

The evidence reviewed in this chapter supports the possibility of personality change, but it does not support unlimited change. Traits are relatively stable for a reason. They are shaped by

many interacting factors, including genetics, biological systems, developmental history, habits, relationships, social roles, and environments. A short intervention can push on some of these factors, but it cannot instantly replace all of them.

There are also limits in measurement. Many personality change studies rely on self-reports. If people know they are trying to become more extraverted or more conscientious, they may later interpret themselves through that goal. This does not mean self-reports are useless, but it does mean we should be skeptical of overly dramatic conclusions. Informant reports, behavioral measures, long follow-up periods, and randomized designs all help, but no single study design solves every problem.

Another issue is durability. A person might show meaningful change during an intervention but drift back toward previous patterns after the intervention ends. This is common in behavior change more generally. A structured environment can make new behaviors easier, but when the structure disappears, old habits and old situations can return. Durable personality change likely requires either continued adherence, new habits that persist without effort, or environments that keep reinforcing the new pattern.

Finally, personality change is not automatically improvement. Higher levels of a socially desirable trait can still have costs in some contexts. More conscientiousness may come with rigidity or stress. More agreeableness may make it harder to set boundaries. More extraversion may be useful in some environments and exhausting in others. Personality change should be evaluated in relation to goals, contexts, tradeoffs, and well-being, not treated as a simple march toward the highest possible score on every desirable trait.

Closing Thoughts About Personality Change

The main conclusion from this chapter is neither ‘personality is fixed’ nor ‘you can become anyone you want by deciding hard enough.’ Both claims are too simple. Personality traits are stable enough that change usually requires sustained effort, repeated practice, and supportive contexts. But they are also malleable enough that interventions, therapy, social roles, and deliberate behavioral practice can shift them.

The state-to-trait perspective gives us a practical framework. If traits describe average patterns of thoughts, feelings, and behaviors, then changing a trait requires changing those patterns repeatedly over time. A desire to change may start the process. A growth mindset may make the process feel more possible. But the crucial step is follow-through: repeatedly enacting the states that are consistent with the desired trait until they become more typical.

This framework should also make us skeptical of magical thinking about personality change. A new label, a motivational quote, or a one-time insight is unlikely to transform a broad trait. Personality change is more mundane than that. It is built from repeated behaviors, recurring situations, feedback from others, and gradual shifts in identity and habit. That

may make personality change sound less dramatic, but it also makes it more scientifically understandable.

References

- Baranski, E., Gardiner, G., Lee, D., & Funder, D. C. (2021). Who in the world is trying to change their personality traits? Volitional personality change among college students in six continents. *Journal of Personality and Social Psychology*, 121(5), 1140.
- Hudson, N. W. (2021). Dynamics and processes in personality change interventions. In *The handbook of personality dynamics and processes* (pp. 1273-1295). Academic Press.
- Hudson, N. W., Briley, D. A., Chopik, W. J., & Derringer, J. (2019). You have to follow through: Attaining behavioral change goals predicts volitional personality change. *Journal of personality and social psychology*, 117(4), 839.
- Hudson, N. W., & Fraley, R. C. (2015). Volitional personality trait change: Can people choose to change their personality traits?. *Journal of personality and social psychology*, 109(3), 490.
- Hudson, N. W., Fraley, R. C., Briley, D. A., & Chopik, W. J. (2021). Your personality does not care whether you believe it can change: Beliefs about whether personality can change do not predict trait change among emerging adults. *European Journal of Personality*, 35(3), 340-357.
- Roberts, B. W., Hill, P. L., & Davis, J. P. (2017). How to change conscientiousness: The sociogenomic trait intervention model. *Personality Disorders: Theory, Research, and Treatment*, 8(3), 199.
- Roberts, B. W., & Jackson, J. J. (2008). Sociogenomic personality psychology. *Journal of personality*, 76(6), 1523-1544.
- Roberts, B. W., Luo, J., Briley, D. A., Chow, P. I., Su, R., & Hill, P. L. (2017). A systematic review of personality trait change through intervention. *Psychological bulletin*, 143(2), 117.
- Stieger, M., Flückiger, C., Rügger, D., Kowatsch, T., Roberts, B. W., & Allemand, M. (2021). Changing personality traits with the help of a digital personality change intervention. *Proceedings of the National Academy of Sciences*, 118(8), e2017548118.
- Tang, T. Z., DeRubeis, R. J., Hollon, S. D., Amsterdam, J., Shelton, R., & Schalet, B. (2009). Personality change during depression treatment: a placebo-controlled trial. *Archives of general psychiatry*, 66(12), 1322-1330.