

Chapter 6

How Stable Are Attachment Styles in Adulthood?

One of the key themes of attachment theory is that people's attachment styles—their typical ways of thinking, feeling, and behaving in close relationships—are relatively stable across time. Indeed, the continuity of attachment styles is what enables theorists to explain a variety of psychologically interesting phenomena, including the tendency for people to recreate maladaptive relationship patterns. But the theory also emphasizes the plasticity of attachment styles. In order for working models to reflect people's relationships experiences, they must be capable of being modified and revised in light of ongoing experiences. Indeed, part of what motivates much of the clinical interest in attachment theory is the assumption that attachment styles can change in adulthood (see chapter: What Are the Implications of Attachment Processes for Psychopathology and Therapy?). If people cannot change, there is little reason to invest in therapy or interventions that might promote security.

The purpose of this chapter is to review what is known about stability and change in attachment styles in adulthood. We begin by reviewing some of Bowlby's ideas concerning stability and change. One of the ideas that Bowlby emphasized was that the processes giving rise to stability and change were often features of the same system. Thus, rather than viewing stability and change as mutually exclusive outcomes of these processes, he viewed them as forces that coexist within dynamic systems. Using these ideas as a general framework, we review research on mechanisms that promote continuity as well as mechanisms that promote change. We then turn to the empirical literature that has attempted to quantify the degree of stability that exists in attachment styles. We also address some of the debates that have ensued over the years about stability and change. We review what is known about the ways in which attachment styles vary across different phases of the adult life course. Finally, we attempt to highlight some of the questions that remain and suggest ways in which researchers can answer those in the future.

We begin by noting that there are many ways to conceptualize and measure continuity in individual differences research (see Caspi & Roberts, 2001, for a review). One of the most crucial distinctions is between what is called rank-order stability and mean-level stability. *Rank-order stability* is concerned with

the ordering of individual differences across time. For example, if we were to assess attachment styles in a sample of individuals at the beginning of the year and then reassess the same sample a year later, we might wish to know whether the people who were highly secure at Time 1 are also the people who are highly secure at Time 2. We could index this quantitatively using a test–retest correlation. If everyone maintained their standing relative to others, then the test–retest correlation would be 1.00. If, however, some of the people who were highly secure at Time 1 are not so secure at Time 2 and others who were insecure at Time 1 are now relatively secure at Time 2, then the correlation would be closer to 0.00. Rank-order stability is one of the most commonly used ways to measure stability in attachment research.

An alternative way to conceptualize and measure stability is mean-level stability. Mean-level stability is concerned with the extent to which the average level of a variable changes across time. If the sample described earlier, for example, had an average score of 5.23 on a 1- to 7-point scale of security at both assessment waves, we would conclude that there was no mean-level change across time. To the extent to which those means are different, however, we may conclude that there are forces leading people to increase (or decrease) in security systematically across time.

Importantly, mean-level and rank-order stability are conceptually and statistically independent (Caspi & Roberts, 2001). It is possible, for example, for the average level of security to remain constant across time, despite there being zero stability in the rank ordering of individuals. This could happen if everyone who is highly secure at Time 1 happens to be highly insecure at Time 2 and vice versa. Similarly, it is possible for mean levels of security to increase (or decrease) across time, even if the rank ordering of individuals is perfectly preserved (ie, test–retest $r = 1.00$). As a result, researchers often address questions about rank-order and mean-level stability separately. This does not mean, however, that the answers to one question are irrelevant for the other; it simply means that one cannot logically conclude that, just because people become less anxious in their attachment patterns across time (eg, Chopik, Edelstein, & Fraley, 2013), the rank ordering of individual differences is changing too. It is quite possible that the same kinds of factors that give rise to mean-level changes in some circumstances also give rise to instability in the rank ordering of individual differences.

With that as context, in this chapter we review research that is primarily concerned with rank-order stability. Specifically, we will focus on processes (eg, breakups) that are assumed to lead to change in attachment at the individual level and, as a consequence, disrupt the rank ordering of individual differences. Near the end of the chapter we will also discuss mean-level changes in attachment, with a focus specifically on how attachment changes as a function of age. We should also note that continuity and change are two sides of the same coin. Although some processes may specifically facilitate continuity in attachment whereas others may facilitate change (see later), the outcomes of both processes

are indexed in the same way (eg, with test–retest correlations). Thus, rank-order coefficients quantify *both* stability and change. The larger those coefficients are, the more stable people are; the closer to zero those coefficients are, the less stable people are.

WHAT DID BOWLBY SAY ABOUT STABILITY AND CHANGE? METAPHORS FOR DYNAMIC PROCESSES

Many of Bowlby's ideas about stability and change in attachment organization were inspired by C. H. Waddington's (1967) discussion of cell development (see Fraley & Brumbaugh, 2004, for an in-depth discussion). Waddington was an esteemed developmental embryologist in Bowlby's time who was trying to understand how a cell may maintain a specific developmental trajectory in the face of varying environmental forces. Waddington and others had observed that, once a cell begins to assume specific functions (eg, it will become part of the visual system), minor changes to the cell's environment are unlikely to alter the cell's developmental trajectory. That is, despite attempts to disrupt the cell's growth, the cell continues as if it has a specific goal in mind. Although a cell has the potential to assume many different functions early in its development, once a specific trajectory has been established, Waddington argued that the trajectory becomes *canalized* or *buffered*, making it increasingly unlikely that the cell will deviate from that developmental course.

To illustrate this process more concretely, Waddington compared cell development to the behavior of a marble rolling down a hill. In Waddington's analogy, the marble represents a cell and the various troughs at the end of the landscape represent alternative developmental functions that the cell can assume. Waddington considered the specific shape of the landscape to be controlled by the complex interactions among numerous genes, and, as such, he referred to it as the *epigenetic landscape* (Fig. 6.1). Once the marble begins its descent, it settles into one of several pathways defined by the valley floors of the landscape. A slight nudge may push the marble away from its course, but the marble will eventually return to the trajectory previously established. As the marble continues along the basin of the specific valley, it becomes increasingly unlikely that external forces will cause it to jump from one valley to the next. Certain features of the marble, such as its momentum, help to keep the marble moving along the existing pathway, and certain features of the landscape itself, such as the steepness and curvature of the valleys, serve to buffer the marble against forces that might disrupt its trajectory.

Waddington's illustration was highly influential in Bowlby's thinking about stability and change in attachment organization. Bowlby noted, for example, that understanding the behavior of the marble in Waddington's example required understanding not only the mechanics of the marble itself, but




FIGURE 6.1 Waddington's epigenetic landscape.

also the environment in which it was situated. Some of the valleys in the hill are much steeper than others and, as a result, are much more likely to steer the marble toward a specific outcome. In the case of human development, Bowlby suggested that certain features of the individual's interpersonal environment can sustain the individual's developmental trajectory. For example, when the individual's family context is stable, he or she is unlikely to experience interactions that challenge his or her representations of the world. The powerful nature of this dynamic was emphasized by Bowlby's (1973) observation that children and adults typically have the same parents, same community, and same broad culture for long periods of time. When the environment is stable, an individual is unlikely to be confronted with experiences that are inconsistent with existing expectations—a point to which we return later in this chapter.

Bowlby also noted that certain properties of the marble itself—its momentum, direction, and speed, for example—also determine its trajectory. In the case of attachment dynamics, Bowlby also called attention to a number of individual-level or psychodynamic processes that may promote continuity. Bowlby observed that people often select environments that are consistent with their preexisting working models. For example, an individual who is relatively secure is more likely, than one who is insecure, to be accepting of others and to interact with them in ways that will help build and establish trust. Moreover, Bowlby argued that basic social-cognitive processes lead individuals

to assimilate new information into existing knowledge structures rather than to build new representations to accommodate for discrepancies.¹

Despite the many forces that promote stability, there are many forces that promote change. In Waddington's metaphor, for example, although a marble is likely to return to the trajectory previously established when it is nudged in one direction or another, a sizable push can knock it into a new valley—especially early in the developmental process. As a consequence, not only does the marble change direction temporarily, its long-term trajectory can be changed as well. Bowlby recognized that there are experiences that people have that have the potential to knock them out of equilibrium, so to speak. The loss of a spouse, for example, can often lead people to question fundamental assumptions they hold about the world (Bowlby, 1980; Fraley & Shaver, 2015), all of which have some bearing on how safe they feel in the world and how responsive they believe people will be to their needs.

One of the important features of the epigenetic landscape metaphor was that it provided Bowlby with a unified framework in which to understand both continuity and change. The framework captured his intuition that, once certain dynamic processes were set in motion, they would be naturally self-sustaining. Indeed, we will discuss some research below on some of the social-cognitive mechanisms that lead working models to reinforce themselves. But the metaphor also provides a framework for understanding how change can occur despite the self-sustaining nature of the system. Specifically, by nudging the marble with enough force or when the marble is at an unstable location in the landscape (eg, near the cusp of a valley), it is possible to modify the marble's trajectory. Indeed, there is a growing body of research that we review below which suggests that certain life events may be capable of creating shifts in people's attachment styles. In the sections that follow we discuss in more depth some of the research on continuity and change in adult attachment styles.

WHAT LEADS TO STABILITY IN ADULT ATTACHMENT STYLES?

According to Bowlby, there are many factors that can facilitate continuity in attachment styles across time. In the sections that follow we discuss two broad mechanisms of stability: those that concern intrapsychic processes and

1. Many of these ideas are commonplace in modern personality and developmental research and theory. For example, Scarr and McCartney (1983), in their classic paper on niche picking, argued that one mechanism that promotes stability involves the way in which people select and construct environments that are compatible with their preexisting, genetically influenced dispositions. Modern research on adult personality development (eg, Caspi & Roberts, 2001; Specht et al., 2014) also emphasizes the ways in which continuity can be maintained through the processes of evocative transactions (ie, other people interact with the individual in ways that sustain his or her dispositions), reactive transactions (ie, people interpreting their environment in ways that are compatible with or biased towards the assumptions they already hold), and proactive transactions (eg, selection effects, such as people choosing social contexts, occupations, or partners that reinforce existing dispositions).

those that concern the structure of the environment. Importantly, each of these mechanisms is thought to be self-sustaining in the way entailed by the epigenetic landscape metaphor. That is, once these processes are set in motion, they have the potential to feed back onto themselves, further sustaining continuity in attachment patterns.

Intrapsychic Processes

Research suggests that the working models people hold may play a role in shaping the ways in which people interpret and understand their interpersonal experiences. The consequence of which is that people can create self-fulfilling prophecies by interpreting the behavior of others in ways that reinforce the assumptions they already have about close relationships.

Attributional Processes and Confirmation Bias

Imagine, for example, that you are at a party with your partner. After mingling a bit and having a few drinks, you notice that your partner is no longer by your side. You wander around a bit, refresh your drink, and, eventually, find your partner in another room. Your partner is talking to someone you do not recognize and is smiling and playfully laughing. How does this make you feel? What do you think their intentions are?

Collins (1996) conducted a study in which people were asked to imagine a variety of scenarios like the one above—situations in which the behavior of a loved one was potentially ambiguous—the behavior could be harmless or could represent a threat to the relationship. Although each participant read identical scenarios, the way participants reacted to the scenarios differed dramatically. Some people believed that their partner was trying to make them feel jealous; other people wrote the event off as it was nothing out of the ordinary. Importantly, Collins (1996) found that how people responded—the attributions they made about their partner’s behavior—was a function of their attachment styles. People who were relatively insecure, for example, were more likely than those who were secure to construe the partner’s ambiguous behavior as a threat to the relationship.

Collins’ research shows, that even when different people are exposed to the same information, the way they interpret that information is biased by their working models. Thus, what people “see” and what they experience tends to reinforce rather than challenge the assumptions they already hold about the world. This dynamic provides one potential mechanism of stability. That is, it is difficult for people to modify their assumptions about the availability and responsiveness of other people in their lives if they are predisposed to view the behavior of others as negligent or insensitive. Vicary and Fraley (2007) expanded on this theme by studying the decisions people made over the course of a narrative about a hypothetical relationship. Specifically, they adopted the

format of a popular book series from the 1970s and 1980s, the *Choose Your Own Adventure* series. In these books, readers assume the role of the protagonist in the story and, at various crucial moments in the narrative, the reader must decide between two or more options (eg, enter the dark cave or turn back and head for safety). Depending on the decision the reader makes, he or she flips ahead in the book to different parts of the story. As a consequence, different choices lead to different narratives, with new choices and different outcomes.

Vicary and Fraley adopted this format to examine the way in which people navigate imaginary relationships. They asked people to imagine themselves in a relationship with a person who behaved in ways that could be ambiguous at times. Participants made 20 choices at various points in the story, choices that had the potential to be beneficial for the relationship (eg, telling your partner that you understand the issues he or she is confiding in you) or destructive (eg, telling your partner that he or she is overreacting). Vicary and Fraley found that, overall, people with insecure attachment styles tended to make poorer choices than people with secure attachment styles (see also Gillath & Shaver, 2007). Although most participants gradually came to make better choices over the course of the interactive story, the rate at which insecure people did so was slower than that for secure people. Stated differently, not only were insecure people biased to transform a potentially ambiguous situation into a negative one, they had a tendency to persist with those detrimental choices, making it difficult for them to recover and steer their interpersonal relations in a more constructive direction.

Zhang and Hazan (2002) argued that confirmation biases can come into play in the kind of information people seek out when evaluating other persons. These authors used a person perception paradigm to investigate the way in which working models may bias how people weigh information about others. Participants were provided with a description of a person (either a classmate or a romantic partner) and were asked to judge how many times the person would have to behave in a way that was inconsistent with the trait in question (eg, considerate) for the participant to be convinced that he or she does not possess that trait. They also asked people how many times the person would have to behave in a way that was consistent with certain traits (eg, trustworthy) to conclude that the person did, in fact, possess that trait. They found that people who were highly avoidant required more evidence to make positive judgments (eg, that the person was considerate or trustworthy) and more evidence to reject a negative judgment (eg, that the person was lazy or rude). These findings suggest that the quantity of behavioral evidence that people use to inform their judgments about others is shaped by their working models. Specifically, people with avoidant attachment styles require a great deal of positive exemplars to disconfirm assumptions about a person's potential negative attributes and require very few negative exemplars to confirm the negative impressions they already hold.

Transference Processes

As people forge new relationships, they sometimes learn that they have recreated the same kinds of relationship dynamics that characterized their previous relationships. Social-cognitive psychologists have attempted to explain this process through the idea of transference (eg, Andersen & Cole, 1990; Andersen, Glassman, Chen, & Cole, 1995). *Transference* is defined as the process by which existing mental representations of significant others are activated and applied to make sense of new social interactions. When interacting with a new person, one observes a number of important cues, such as whether the person is a smoker, the color of his or her hair, whether the person is outgoing or shy, and so on. The presence of these cues has the potential to activate mental representations of other people who possess similar attributes. As a result, representations of significant others have the potential to color the way we perceive new individuals and, in some cases, “go beyond the information given” when attempting to make inferences about the new person.

In a prototypical transference study, people visit the lab and provide information about significant others from their lives. Then, in a separate and “unrelated” study weeks later, people may interact with a real or hypothetical individual. For participants in the experimental condition, the new person is designed to have certain features that are similar (but not necessarily identical) to those of the person’s significant other. For example, if people described the significant other as a poet, the new person may be described as a writer. In a yoked control group, different people learn about the new person, but, for them, the individual does not have features that resemble a significant other from their past. After learning about the new person, all participants are then asked to perform a new task. They may be asked, for example, to write down as much as they can remember about the new person they read about. The outcome of interest in such a study is whether people “remember” things about the new person that were not actually present, but were true, in fact, of the person’s significant other.

Brumbaugh and Fraley (2006) used this paradigm to study the ways in which mental representations of significant others may guide the way in which individuals relate to novel people. Specifically, in the first session, participants were asked to nominate a significant other from their past and provide a number of facts about the person, including the person’s traits, interests, and habits. Then, 2 weeks later in an unrelated study, people came to the lab to view and evaluate personal ads on a website. In the experimental condition, the ads were constructed to contain a few features that were similar to those of the person’s significant other. Participants in the yoked control condition saw the same ad, but it was not based on their significant other. After viewing the ads, participants rated how secure or insecure they thought they would feel in a relationship with the person.

Brumbaugh and Fraley found that people who were insecure with their significant other were also more likely to feel insecure when imagining what it

would be like to be dating the person described in the ad. This was true regardless of whether the ad contained information that was similar to that of the significant other. Stated differently, people who were insecure in past relationships were likely to feel insecure with respect to a new potential partner. But, beyond that general association, Brumbaugh and Fraley (2006) found that this effect was magnified when the ad was designed to resemble the subject's significant other. That is, subtly activating the significant other representation led people to relate to the person described in the personal ad in a way that was congruent with preexisting patterns of attachment.

One of the noteworthy findings in the Brumbaugh and Fraley study was that, although people were more likely to feel insecure about the people described in the personal ads if an insecure significant other representation had been primed, they nonetheless expressed a greater interest in dating the person described in the ad. Thus, it appears that, at least in some cases, a feeling of familiarity can trump an assessment of security when deciding whether a potential partner is a desirable option.

Selection and Attraction Processes

Chappell and Davis (1998) proposed the attachment-security hypothesis; the idea that, when given a choice between partners who are potentially secure and partners who are potentially insecure, most people will choose the secure partner as more desirable. According to Chappell and Davis, this is expected because the attachment behavioral system is designed to seek cues that others are responsive and available; attributes that are more typical of secure than insecure partners. One consequence of this attraction process is that insecure people may seek out potentially secure partners when developing new relationships. Indeed, in their early research Chappell and Davis found that, when people were given a choice between hypothetical partners who varied in their attachment security, people tended to choose the secure-seeming partner over those who were designed to be insecure.

Despite the tendency for most people to find secure prospects more attractive, not everyone ends up in relationship with a secure individual (we review some of this work in more depth in chapter: What Are the Effects of Context on Attachment?). Part of the explanation is likely due to market forces: there are not enough secure people to go around. But part of the explanation may also have to do with the ways in which insecurities play out in relationship contexts. Namely, highly anxious people, for example, might be viewed as bad relationship partners, making secure-insecure pairings relatively unstable (eg, Pietromonaco & Carnelley, 1994). Moreover, insecure people may drive secure partners away in dating contexts. In a striking demonstration of this process, McClure and Lydon (2014) studied people in a speed-dating paradigm and found that individuals who were more anxious with respect to attachment were more likely to come across in undesirable ways, expressing greater verbal disfluencies and interpersonal awkwardness. These interpersonal behaviors, in

turn, have the potential to undermine the formation of intimate relationships, potentially reinforcing the insecurities that highly anxious people already have.

Although both highly secure and insecure people tend to prefer partners who are secure, people tend to end up with others who are similar to them with respect to attachment (Holmes & Johnson, 2009). People who are relatively anxious, for example, tend to be paired with others who are relatively anxious. Thus, it seems that something takes place in the development of romantic relationships that leads insecure people to be with others who are insecure. It is unclear whether couples who are discordant with respect to attachment styles are more likely to breakup early in the relationship formation process or whether, in the process of mutually influencing one another, partners become more alike in the way they think, feel, and behave with respect to attachment-related concerns (see further for more on this possibility).

Environmental and Relational Processes

The processes we have discussed up to this point emphasize what takes place within the individual—social-cognitive mechanisms that lead people to see what they expect and to recreate interpersonal patterns with which they are familiar. But another source of stability emerges from the structure of the person's interpersonal world and the way in which the person shapes it and the way it shapes him or her. For example, if two people are in a relationship, they are in a position to mutually influence one another. Thus, not only is one person interpreting the behavior of the other in a way that is consistent with his or her own working models, but the partner is also doing the same. The consequence of this is that the couple is essentially engaged in a dyadic process in which they are reinforcing one another's attachment patterns; they function like atoms that are orbiting around one another in a dynamically stable pattern.

Hudson, Fraley, Brumbaugh, and Vicary (2014) examined this process in a longitudinal study of couples who were followed five times over the course of a year. Specifically, these researchers examined the way in which each individual in the couple, related to the other (ie, partner-specific attachment styles; see chapter: What Are Attachment Working Models?) and studied the way in which those attachment styles changed jointly across time. They found that, in general, people did not show a strong tendency to systematically increase or decrease in attachment anxiety or avoidance across the year. But people's attachment styles did vary across time; the same person, for example, was more secure on some occasions than others. Importantly, these person-specific deviations tended to be correlated within couples. That is, on occasions on which a person was feeling more insecure than usual, his or her partner was also likely to feel more insecure than he or she would typically feel (see also Davila, Karney, & Bradbury, 1999). Thus, the idiosyncratic changes that people experienced tended to be shared with, and in some cases, possibly even influenced by, the partner (see Sbarra & Hazan, 2008, for an in-depth discussion of these kinds of

processes). This suggests that one source of stability in adult attachment styles stems from dyadic processes. Simply being in a close relationship with another person creates a system of mutual influence that leads people to converge to some degree in their attachment styles and reinforce that pattern of relating. Thus, if a relatively secure person begins to drift in a more insecure direction, the other partner may pull him or her back.

WHAT LEADS TO CHANGE IN ADULT ATTACHMENT STYLES?

According to Bowlby (1973), people tend to assimilate ongoing experiences into the working models that they already have. Thus, when a partner behaves in a way that is slightly at odds with one's existing expectations, one is more likely to perceive the interaction as being consistent rather than inconsistent with one's expectations (Collins, 1996). But Bowlby (1973) also argued that working models are responsive to ongoing relational experiences. Thus, if one's experiences sufficiently challenge one's existing expectations, those experiences have the potential to lead to changes in attachment organization (Fraley, 2002).

What kinds of factors lead to change in adult attachment styles? In the sections below we review some of the events and experiences that have received the most attention in the attachment literature. This review is not meant to be exhaustive, or to imply that the absence of certain factors (eg, the loss of a loved one) from the review implies that such factors may not be relevant to understanding change. We should also note that, in many cases, these factors represent the flip-side of processes that may facilitate stability. For example, to the extent to which relationship breakups may lead to instability in attachment style, the persistence of a committed relationship represents a factor that promotes continuity in attachment.

Major Life Transitions

A large body of research has investigated the implications of major life transitions, such as the transition to parenthood (eg, Simpson, Rholes, Campbell, Tran, & Wilson, 2003) or the transition to college (Lopez & Gormley, 2002), for understanding attachment dynamics. Although work in this area has been broad—examining a variety of issues beyond questions of stability and change—the work is clearly relevant to basic questions about change, such as whether people undergoing major transitions are more or less likely to exhibit stability in their attachment styles.

Why the focus on major life transitions? Simpson and colleagues argue that lawful change is most likely to occur when individuals face a stressful, life-altering event (Simpson, Rholes, Campbell, & Wilson, 2003). The reason for this is that such events expose people to new information and experiences that create opportunities for one's core assumptions to be challenged. Moreover, such experiences may lead people to reflect upon or reevaluate the assumptions

they hold about themselves, their partners, and their relationships. In the sections below we discuss two major life transitions—parenthood and breakups—that have been studied extensively in the attachment literature.

Parenthood

The transition to parenthood is one major life transition that has received a lot of attention in adult attachment research. Having a child can be a stressful experience, one that has the potential to tax people's interpersonal resources considerably. Moreover, the birth of a child has the potential to rekindle significant attachment-related experiences from the expecting parent's past, leading the individual to reflect upon his or her own developmental experiences and to consider ways in which one may wish to parent differently. Becoming a parent has the potential to lead to other social-structural changes as one begins to socialize with other parents with same-age children and interact with local educational communities and teachers. Parenting can also be stressful for couples as they find themselves with less time for adult activities that they may have enjoyed previously (eg, dining out, theater) and potentially struggling with the negotiation of child care responsibilities.

Feeney, Alexander, Noller, and Hohaus (2003) examined the association between adult attachment styles and depression during the transition to parenthood. Specifically, they assessed 76 couples who completed surveys during the second trimester of pregnancy, and 6 weeks and 6 months after childbirth. Importantly, Feeney and colleagues also studied an age-matched control sample of 74 childless couples, thereby allowing them to draw comparisons between attachment processes for couples who were and who were not undergoing the transition to parenthood. Feeney and colleagues found that attachment-related anxiety was less stable for wives undergoing the transition (test–retest correlation of approximately 0.54) than for other participants (test–retest of approximately 0.72). (In addition, attachment anxiety predicted increases in new mothers' depressive symptoms across time. Women who were more insecure preterm, in other words, were more likely to experience symptoms of depression after the birth of their children.)

Simpson et al. (2003a) examined a sample of approximately 100 couples both 6 weeks before and 6 months after childbirth. Importantly, they found that, on average, attachment styles did not change from pre- to postbirth. Thus, the transition to parenthood per se did not lead people to become more (or less) secure. However, there were individual differences in the extent to which women sought support from their spouses in the prenatal assessment. They found that women became more anxious with respect to attachment if they entered into parenthood perceiving less support from their spouses and more spousal anger. Moreover, women who entered parenthood seeking less spousal support became more avoidant in their attachment across the transition. Thus, although the transition to parenthood did not lead to mean-level changes in attachment

security across time, the way in which people navigated the transition was consistent with—and potentially sustained—their attachment styles.

Breakups

Another significant life transition is relationship dissolution. The breakup of a romantic relationship has the potential to lead to substantial disruption in attachment processes (Sbarra & Hazan, 2008). For example, if an exclusive relationship ends because one of the individuals has been unfaithful, this experience is likely to shatter the sense of trust that exists between partners. This may have implications for the extent to which the person feels that he or she can open up to or depend on others and, potentially, may make it more challenging for the person to fully trust his or her next partner.

A few empirical studies have examined the potential impact of relationship dissolution on attachment styles. In one of the classic studies on this topic, Kirkpatrick and Hazan (1994) followed a sample of individuals over a 4-year period. They found that changes in relationship status were associated with changes in attachment. Specifically, of the participants who were involved in a romantic relationship at the initial assessment, 90% of secure individuals who did not experience a breakup were secure 4 years later whereas approximately 50% of those who experienced a breakup remained secure. These findings suggest that relationship dissolution has the potential to undermine the sense of security that people feel in close relationships.

Ruvolo, Fabian, and Ruvolo (2001) examined a sample of 301 dating couples longitudinally and found that women became less secure after a breakup. They also found that people were likely to become more secure across time if they were involved with the same partner over the course of the longitudinal study. Scharfe and Cole (2006) examined stability and change in attachment among a sample of university students who were graduating. They found that relationship status partially moderated the stability of attachment, such that the test–retest stability of attachment was lower among those who changed their status compared to those who did not (ie, those who stayed single or stayed coupled).

In summary, it appears that the loss of a romantic partner has the potential to undermine people's security. To be clear, however, these associations are relatively weak. This is probably because most of the research to date has not been able to carefully evaluate whether the relationship was ending for reasons that were agreeable to both individuals. In addition, breakups themselves are not always discrete events; they represent a seemingly arbitrarily timed transition in a relationship that is already unsatisfying or not mutually rewarding. Thus, any change that has taken place in attachment styles is likely to have occurred before the breakup itself took place. Finally, not all breakups are as tumultuous as we sometimes assume in our culture. The end of a dissatisfying relationship can be a positive event for some people and, as a result, may not challenge their working models in ways that facilitate dramatic change. Thus, although

research suggests that breakups facilitate attachment change, we would not want to assume that the effects are dramatic.

War-Related Trauma

Mikulincer, Ein-Dor, Solomon, and Shaver (2011) assessed the 17-year trajectories of attachment orientations in two groups of Israeli veterans from the 1974 Yom Kippur war. One group was comprised of ex-prisoners of war and the other was a comparison group of veterans who had not been held captive. Both groups of veterans completed measures of adult attachment styles at 18, 30, and 35 years after the war, along with a variety of other measures, including PTSD symptoms. Mikulincer and colleagues found that, overall, ex-prisoners of war were less secure than those in the comparison group and that, while the comparison group generally became more secure across time, ex-prisoners of war became more insecure across time. They also found that the experience of PTSD at each assessment heightened feelings of insecurity at that time point, beyond the potential effects of other variables.

These findings are important for at least two reasons. First, they reveal that specific events have the potential to lead to long-term changes in attachment style. Second, not only were these veterans less secure than others 18 years after the war, they were on a *trajectory* toward greater insecurity across time. Thus, the experience of being a prisoner of war led not only to shifts in security, but shifts in the developmental time course of security. In Mikulincer's words, this traumatic experience had a "long-term pathogenic effect."

Relationship Conflict and Support

According to attachment theory, the security that a person experiences at any one moment is derived from the knowledge that others are available and accessible. As a result, people should be more likely to feel secure in their relationships on occasions in which they perceive their partner as being supportive and responsive. Similarly, they should also feel insecure on occasions when their partners are not supportive—occasions when they perceive conflict, excessive distance, or a lack of mutual understanding and respect.

There are now a number of intensive longitudinal studies, which suggest that fluctuations in security hinge on the state of interpersonal relationships. For example, Holman, Galbraith, Timmons, Steed, and Tobler (2009) assessed the extent to which individuals perceived threats to the availability of their parents and their romantic partners. Specifically, they assessed the extent to which people felt that their attachment figures were accessible, responsive, and openly communicative. Holman and colleagues found that people who perceived greater threats to the availability of their attachment figures were less likely to be secure 1 year later.

Green, Furrer, and McAllister (2011) examined stability and change in attachment style in a sample of 181 low-income mothers. The mothers were

assessed across three waves, starting shortly after the birth of their children. Green and colleagues found that increases in social support led to decreases in attachment-related anxiety. And, although decreases in anxiety did not lead to prospective increases in social support, decreases in attachment-related avoidance did.

La Guardia, Ryan, Couchman, and Deci (2000) examined attachment security in a 30-day daily diary study. On days in which people felt that their basic needs for autonomy, competence, and relatedness were met, they experienced greater attachment security. (And, when those needs were not met, they felt more insecure relative to their average.)

Chow, Ruhl, and Buhrmester (2014) examined a sample of approximately 300 adolescents from 6th to 12th grade. They found that attachment-related avoidance prospectively predicted friendship exclusion. In turn, friendship exclusion prospectively related to avoidant attachment. One of the valuable features of this study is that it demonstrates that there may be bidirectional influences between attachment security and relational experiences in adolescent friendships. That is, it seems that not only do friendship experiences have the potential to change attachment, attachment has the potential to shape friendship experiences.

The Meaning of Life Events

In their review of the literature on change, Davila and Sargent (2003) noted that many studies failed to provide convincing evidence that specific life events were related to change in attachment. They suggested that one reason why it has been so challenging for researchers to demonstrate consistently that specific life events, such as relationship breakup, might be associated with changes in attachment style is that these events do not have the same meaning for everyone. Some people who end a relationship, for example, may be relinquishing ties to someone who was making them miserable. Others, in contrast, may be truly heartbroken and the experiences of the breakup may prompt them to revise their working models of close relationships. According to Davila and Sargent it is necessary to understand the way specific life events are construed in order to know their implications for attachment change.

To address this issue, Davila and Sargent (2003) studied a sample of approximately 150 students who were asked to complete measures of their attachment style every day for 56 days. In addition, students were asked to indicate whether specific life events (eg, taking an exam, having a fight with a partner) had occurred each day. Importantly, Davila and Sargent asked people to indicate the extent to which each event that had occurred was indicative of an interpersonal loss as indicated, for example, by ratings of the extent to which the event led to a loss in emotional support, friendship, or trust. Thus, in addition to having information about whether or not specific events had taken place, Davila and Sargent also had information on that event's meaning to the participant.

Davila and Sargent found that, in general, the kinds of events most people would construe as negative did, in fact, predict decreases in security across time. When Davila and Sargent also included the meaning of those events in their analyses, however, the predictive value of the events themselves was weakened considerably. Instead, the meaning of the events was what was most predictive of attachment-related change. Specifically, on days on which people experienced events that they construed as being interpersonal losses, they were more likely to experience increases in attachment-related anxiety. Importantly, Davila and Sargent also assessed academic-related losses and stressors and found that those were not related to changes in attachment. Instead, change in attachment style was associated uniquely with changes in interpersonal rather than academic experiences.

Zhang (2009) conducted a follow-up study in which 30 individuals completed measures of daily events twice a week for 4 weeks. Zhang specifically targeted the occurrence of relatively mundane events (rather than low base-rate events, such as breakups) that might be relevant to relationship functioning. Zhang found that when people experience negative daily events, they tend to experience increases in attachment-related anxiety relative to their own baseline. Moreover, when people were asked to rate the implications of these events for interpersonal loss, the extent to which the events were loss-related further contributed to change.

Stable Vulnerability Factors

Davila and colleagues have also advanced the idea that some individuals are more inclined to experience changes in their attachment orientation than others due to stable vulnerability factors. Specifically, individuals who have a history of depression and psychopathology—in their family or personally—are going to have an identity that is less stable than that of others. As a result, they are less likely to have a consistent attachment style across time. Davila, Burge, and Hammen (1997) assessed attachment styles in 155 women several times over the course of a 2-year longitudinal study. They quantified change in a few ways, such as the standard deviation in a person's attachment scores across time. People with high deviation scores experienced more change than those with smaller deviation scores. Davila and colleagues found that those who experienced more change were also likely to report a history of psychopathology and personality disturbance.

Therapy

There have been a wide variety of interventions inspired by attachment theory over the past 20 years (see chapter: *What Are the Implications of Attachment Processes for Psychopathology and Therapy?* for an in-depth discussion of these issues). Although the objective of these interventions varies (eg, to improve marital function and communication), a reasonable question to ask is whether

they lead to changes in people's attachment styles. Taylor, Rietzschel, Danquah, and Berry (2015) reviewed the research literature on this issue, examining both studies that used self-reports and interview-based methods for assessing attachment. They found that, overall, attachment security tends to increase following therapy. Specifically, attachment-related anxiety tends to decrease, but the findings were less clear on whether avoidance also changed. Moreover, the findings seemed to be largely consistent across different methods of assessing attachment (eg, interviews, self-reports) and different patient and therapy variants (eg, therapeutic approaches, settings, and patient groups).

HOW STABLE ARE ATTACHMENT STYLES IN ADULTHOOD?

The research reviewed up to this point indicates that there is evidence for mechanisms that promote stability as well as mechanisms that facilitate change. But how do these different processes stack up? Do the various forces promoting continuity overshadow those promoting change? Do these two kinds of processes balance one another out?

One of the long-standing debates in the study of adult attachment concerns the stability of attachment styles. As Baldwin and Fehr (1995) observed, one reason attachment theory was appealing to many psychologists was that it suggested that attachment styles were trait-like or dispositional variables. That is, it was assumed that a person could be characterized as having a single attachment style and that this attachment style captured the person's patterns of behavior and emotion across both time and circumstance. Baldwin and Fehr (1995) were two of the first researchers to call this assumption into question empirically. They surveyed the literature and found that, in fact, approximately 30% of people tend to report a different attachment style when surveyed on more than one occasion. Thus, if a person reported having an avoidant attachment style at, say, the beginning of the semester, there was a 30% chance that the same person would report a different attachment style a few weeks later. This degree of instability is clearly problematic if one assumes that attachment styles should be relatively trait-like across time.

Baldwin and colleagues also observed that, when people are asked to think of relationships in which they felt relatively secure, avoidant, or anxious, most people are able to do so. That is, the same person can call to mind relationships in which he or she felt secure as well as relationships in which he or she felt insecure. This indicates that a single attachment style does not capture all of the important relationships that a person has. Some people, for example, may be relatively secure in their relationships with their mothers, but less secure in their relationships with their fathers. Moreover, Baldwin and colleagues showed that it is possible to manipulate the attachment style that people report simply by having them bring to mind different kinds of interpersonal experiences (see chapter: What can Social Cognition and Priming Tell us About Attachment?). Thus, it is possible, at least momentarily, to make someone who would

otherwise report being secure report being insecure simply by making certain experiences from their past more salient.

As an alternative to the dispositional view, Baldwin and Fehr (1995) advanced a social-cognitive model of attachment styles. Specifically, they argued that most people have working models that are consistent with multiple attachment styles, but that some of those models are more available and accessible than other ones (see chapter: What Are Attachment Working Models?). Thus, someone who self-reports being secure in his or her relationship may do so not because he or she has a secure disposition, but because secure working models are more likely to be accessible to him or her at the moment of assessment.

Can Some of the Observed Instability be Understood as Measurement Error?

Baldwin and Fehr (1995) concluded that the measurement of attachment, while not being perfect, was good enough and was unlikely to explain the high degree of instability that they had reported. Other researchers challenged this claim, however. Most of the work Baldwin and Fehr (1995) summarized was based on categorical models of attachment (see chapter: How Are Individual Differences in Attachment Measured?). Fraley and Waller (1998) argued that, if attachment styles are not truly categorical, then a substantial amount of instability will exist simply due to cases that are near the threshold. Specifically, if the majority of people are in the middle of the two-dimensional space, minor movements in that latent two-dimensional space will lead to dramatically different attachment classifications, making a secure person at one point in time seem preoccupied at another point in time. Indeed, Fraley and Waller (1998) demonstrated that, in a simple two-dimensional situation in which the variables are normally distributed and there is no true change in attachment, the imprecision of categorical systems will nonetheless lead to observations of 30% of people “changing” their attachment styles across time. Thus, on the basis of measurement error alone, it is possible to explain why approximately 30% of people report different attachment styles across time when categorical assessments are used.

Scharfe and Bartholomew (1994) were the first team to attempt to estimate the continuity of attachment style using methods that (a) did not rely exclusively on categorical assessments and (b) were designed to factor out random measurement error as much as possible. Using latent variable modeling methods, they found that the estimated continuity of attachment over an 8-month period was equivalent to a test–retest correlation of approximately 0.80. Sibley, Fischer, and Liu (2005) found a similar result over a period of 3 weeks. Specifically, about 85% of the variance in attachment measures was shared across measurement occasions (equivalent to a test–retest correlation of approximately 0.92).

What do these recent findings mean? First, they suggest that, for the most part, attachment styles, when measured continuously and over intervals ranging from a few weeks to a few months, are highly stable in adulthood. Although

there is clearly some change taking place (see our review in previous section), one could become reasonably wealthy by betting that people who are highly secure today will also be highly secure months from now. Second, these findings suggest that one cannot easily conclude that attachment is not a “general disposition or trait” (Baldwin & Fehr, 1995, p. 247) on the basis of lack of stability when the measurements themselves are highly imprecise. When better measures are used and measurement imprecision is taken into consideration, estimates of stability increase.

What Does it Mean to say that Something is a General Disposition or Trait? Trait-State Models of Adult Attachment

One of the potential complications underlying debates about continuity and change is that dispositional and social-cognitive perspectives are sometimes treated as if they are mutually exclusive. Namely, many scholars assume that, if something is a trait, it does not require a social-cognitive explanation (eg, Costa & McCrae, 1994). And, conversely, that if something has a social-cognitive explanation, then it is not a trait or disposition (eg, Baldwin & Fehr, 1995; Cervone & Shoda, 1999; Reynolds & Branscombe, 2014).

These stances, however, do not consider the possibility that attachment can have both trait-like and state-like properties. Take room temperature as an example. At any one moment in time a room has a measurable ambient temperature; it may feel too cool, too warm, or just right. However, that temperature can easily change if, for example, the window is opened or if a heater is turned on. In such cases, the temperature of the room now might be quite different from the temperature 60 min ago. Nonetheless, many people would acknowledge that rooms tend to have dispositional qualities with respect to temperature. We complain about some rooms being too drafty, others being too hot, whereas other rooms tend to not attract our notice at all. Due to a combination of factors (eg, quality of the insulation, the efficiency of the heater, number of exterior walls, height from the ground), the typical temperature of various rooms tends to be relatively consistent and predictable. This predictability does not imply that the temperature of a room cannot change suddenly and dramatically. Nor does the ease with which the temperature can be adjusted imply that there are not relatively stable differences between rooms across time.

In short, something as familiar as room temperature can be easily construed as having state- and trait-like properties. The same is the case for attachment. Fraley and colleagues have argued that, in theory, attachment styles can be viewed as being trait-like *and* state-like (Fraley, 2002; Fraley & Roberts, 2005). That is, it should be possible for attachment styles to reflect on-going experiences in the way conceptualized by Baldwin and Fehr (1995). If one has an argument with one’s romantic partner, for example, that conflict is likely to undermine the sense that one feels understood and accepted by that person. This might not make a highly secure person swing all the way to the depths of

insecurity, but it is likely to produce some degree of change. But Fraley and colleagues also argued that the change is likely to be state-like or temporary. Unless the conflict persists, the person will gradually revert to his or her prior levels of security. To refer back to Waddington's metaphor, a small perturbation is unlikely to affect the long-term trajectory of the marble as it rolls down a hill. But that does not mean that the perturbation is unimportant or inconsequential.

These ideas can be formally modeled using variations of trait-state models (Fraley, 2002; Fraley & Roberts, 2005; Kenny & Zautra, 2001). According to one such variation, which Fraley (2002) refers to as a *prototype model* (see also chapter: How Do Individual Differences in Attachment Develop?), a person's attachment security at any one point in time is a function of a stable value of security (ie, a trait-level), previous levels of security (ie, an autoregressive component), and state-like factors that might lead to deviations in security (eg, responsive or conflictual interactions with an attachment figure that may lead to changes in security). These various components are illustrated in Fig. 6.2. One of the valuable features of using a formal model is that one can modify its parameters and study the consequences of those modifications for the predicted test–retest correlations. Experimenting with the parameters of the trait-state model reveals a few important things. First, the model predicts that the overall degree of stability, expressed as a test–retest correlation, observed between attachment measured across two time points can be large or small—even if there is a stable, trait-like source of variance in attachment styles. This is an important observation because it is often assumed that low stability is incompatible with trait models (eg, Baldwin & Fehr, 1995). And, similarly, it is sometimes assumed that high levels of stability

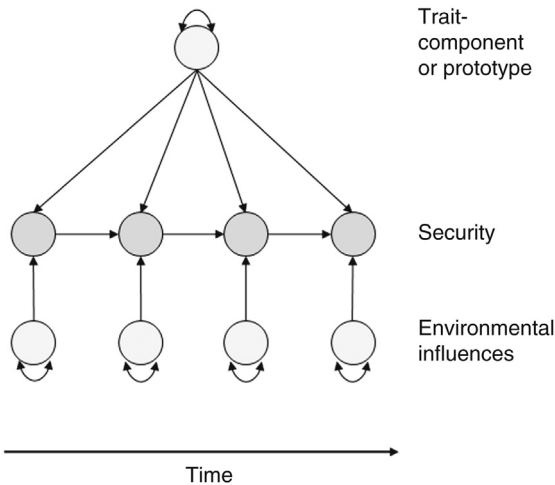


FIGURE 6.2 The prototype model of adult attachment, a trait-state model of continuity and change. The model assumes that variation in security at different assessment waves is a function of (1) a stable trait component, (2) autoregressive processes (ie, security at any point in time t is a function of itself at $t - 1$), and (3) environmental influences.




FIGURE 6.3 Test–retest correlations for security as predicted under different parameter values by a trait-state model. (A) Illustrates predicted correlations for a state-only model under various assumptions. As highlighted by the gray bar, the test–retest correlations over a 5-week interval can be as high (eg, 0.80) or low (eg, 0.10) depending on the specific parameter values assumed. (B) Illustrates predicted correlations for a trait-state model under various assumptions. As highlighted by the gray bar, the test–retest correlations over a 5-week interval can be as high (eg, 0.80) or low (eg, 0.10) depending on the specific parameter values assumed. Overall, these figures reveal that the magnitude of a test–retest correlation across an arbitrary test–retest interval can be high or low in models that assume stable traits/dispositions as well as in models that do not. But they also reveal that the patterns of those associations across time are distinctive. A state-only model predicts decaying test–retest stabilities across increasing intervals; a trait-state model implies that those associations will approach a nonzero value in the limit.

are incompatible with state models. When trait-like assumptions are formalized, however, it can be shown that a trait model can account for low or high levels of stability (Fraley, 2002; Fraley & Roberts, 2005).

Fig. 6.3 illustrates this finding more clearly. The left-hand panel illustrates the way in which the test–retest correlations between two measurements of attachment style vary across increasing test–retest intervals. Fig. 6.3A illustrates the expected test–retest correlations under the assumptions of a model that is a pure state model (ie, the trait variance is set to 0). Notice that the model is capable of predicting both low and high test–retest correlations across a 5-week test–retest interval, depending on some of the parameters of the model. The upper curve in Fig. 6.3A, for example, shows an example in which the test–retest correlation over 5 weeks is as high as 0.80; the lower curve shows an example in which the test–retest correlation over the same period of time is 0.10. Fig. 6.3B illustrates the expected test–retest correlations under the assumptions of a model that assumes both trait-like and state-like processes. Notice that this model is also capable of predicting both low and high test–retest correlations across a 5-week test–retest interval depending on the parameters in the model. Thus, even when there is a stable trait giving rise to security, it is possible for the

test–retest correlation to be quite small. In short, although it is obviously useful to know the magnitude of stability that exists between any two time points, this information is not useful for determining whether attachment is trait-like (ie, dispositional) or not.

Does this mean that empirical data cannot be used to examine the extent to which individual differences can be understood as resulting from trait-like and state-like processes? Not necessarily. The second implication that follows from this model is that the way to distinguish trait-state models from nontrait models (eg, strict contextual or social-cognitive models) lies in the *patterns* (not the magnitude) they predict in test–retest correlations across time. This can be seen by comparing the curves illustrated in the two panels of Fig. 6.3. Although trait models are not precise enough to suggest whether stability will be high or low in an absolute sense (see earlier section), trait models do make a risky prediction, namely, that the degree of stability, whether high or low, will not get increasingly smaller across increasing assessment intervals (Fig. 6.3B). At some point the test–retest stability will stabilize at a nonzero value. This suggests that it is possible to forecast what people will be like in the future with the same degree of precision over 20 weeks as 10 weeks. In contrast, a state-only model predicts that the stability of individual differences will get smaller and smaller as the interval between assessments increases, approaching zero in the limit (Fig. 6.3A). This does not mean that there is *no* stability across assessment waves. But it does imply that one’s ability to predict individual differences in the future will become less accurate the further into the future the prediction is made. The association between measurements of security across 2 months might be quite high, depending on the parameter values in question. But the correlation between assessments of security across 2 years will be smaller (see the solid curve in Fig. 6.3A).

The third implication revealed by these simulations is that there is an asymmetry in stability when one is moving forward versus backward in time (Fig. 6.4).




FIGURE 6.4 The predicted (under hypothetical parameter values) test–retest correlations between measures of security at wave 1 and each subsequent wave (A), wave 15 and all waves prior to and following it (B), and wave 30 and each previous wave (C) under the assumptions of a trait-state or prototype model. These graphs illustrate the asymmetry in predictions about stability that can emerge when moving forward through time versus looking backward through time.

Stated in a way that sounds less like the premise of a science fiction novel: The model predicts that the overall amount of stability observed among *adults* over a constant time period (eg, 1 year) will be higher than the amount of stability observed among *children* over the same time period. As a result, it is easier to know where someone is going than to know where they have been.

One can see this more clearly by examining the middle panel of Fig. 6.4. This figure illustrates the expected test–retest correlations between security measured at wave 15 and all the assessments that precede or follow it. Notice that the predicted test–retest correlation between wave 15 and wave 1 is lower ($r = 0.29$) than the predicted test–retest correlation between wave 15 and wave 30 ($r = 0.44$). The implications of this point are profound because it indicates that, although it might be challenging to accurately infer whether a secure adult was also secure as a child, it is a comparatively safe bet to infer that a secure adult will continue to be secure in the future. This particular prediction dovetails nicely with Bowlby’s use of Waddington’s metaphor for canalization. Essentially the model implies that, as time progresses, people become increasingly entrenched in the ways in which they relate to others. The consequence is that they exhibit greater stability later in life than they do early on. A model that does not posit a role for trait/dispositional factors does not make this prediction.

We should note that this latter prediction regarding asymmetry is not restricted to childhood versus adulthood in practice. Formally, the prediction emerges because in the initial phases of the process, security is a function of the latent trait, but later in the process it is both a direct *and* indirect function of the latent trait. The consequence of this is that the asymmetry can emerge in any system in which there is a plausible “starting point,” such as in the developmental case where children are born into families or when a couple begins dating. In other words, canalization processes emerge naturally within this model and apply to any situation where there is a natural beginning to the dynamic process.

It is difficult to test the various predictions of the prototype model without assessing attachment styles across multiple waves; the traditional two-wave, test–retest design cannot speak to the issues. Fortunately, there are datasets available now that enable these various predictions to be evaluated. Fraley, Vicary, Brumbaugh, and Roisman (2011) reported on data from two samples that were designed to evaluate trait-state models of attachment. The first sample was assessed once a day for 30 days. The second sample was assessed once a week for up to a year. Importantly, attachment was assessed in a contextual fashion (see chapter: How Are Individual Differences in Attachment Measured?). That is, attachment was assessed in relation to people’s parents and their romantic partners separately.

One of the important findings in this research was that the data were more consistent with a trait-state model than a model that did not assume a trait-like source of variance (ie, a pure contextual or social-cognitive model). Thus, although people clearly changed their attachment styles across time, those changes did not accumulate. As a result, the degree of instability observed did

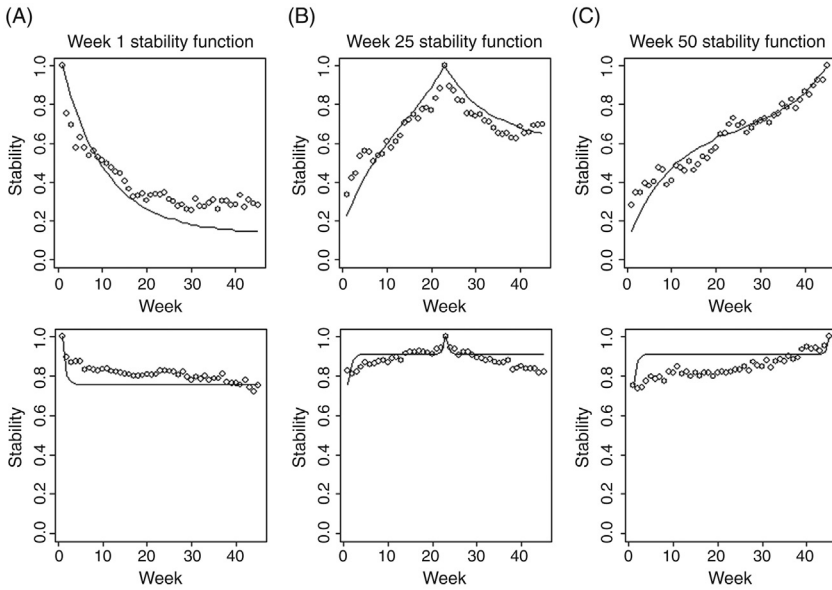


FIGURE 6.5 Empirical test–retest correlations between measures of security at wave 1 and each subsequent wave (A), wave 25 and all waves prior to and following it (B), and wave 50 and each previous wave (C) from Fraley et al. (2011b). The top row shows the data for partner-specific avoidance; the bottom row shows the data for avoidance with mother.

not get lower and lower as the delay between assessments increased. To illustrate, the results for attachment-related avoidance, measured in the context of the relationship to romantic partners, are illustrated in the top row of Fig. 6.5. As can be seen, although the test–retest stability initially gets lower as the test–retest interval increases, it eventually levels off around a value of approximately 0.35 and stays there. As a result, the degree of stability observed over 20 weeks is comparable to that observed over 40 weeks. The second row illustrates the observed correlations for attachment-avoidance with respect to people’s relationships with their mothers. Again, the associations tend to exhibit stabilizing properties; the test–retest correlation between initial measurements of avoidance and avoidance measured 20 weeks later is comparable to that between initial avoidance and avoidance measured 40 weeks later. Taken together, these findings provide strong evidence that there is something dispositional about attachment styles, something that remains invariant despite short-term fluctuations in security.

The second and third panels in the first row of Fig. 6.5 highlight the asymmetry in continuity across time. Namely, the test–retest correlation between partner-avoidance measured in week 23 and week 13 (a 10-week delay) was lower than the correlation between partner-avoidance in week 23 and week 33 (also a 10-week delay). This is compatible with the canalization effects emphasized by Bowlby (1973) and implied by the trait-state model.

Does the Stability of Attachment Vary Across Relationship Types?

In chapter: What Are Attachment Working Models?, we highlighted the idea that attachment styles can vary in their degree of specificity. That is, theoretically, people have general attachment styles, but they also have unique patterns of attachment that characterize the way they relate to specific people in their lives, such as their parents or their romantic partners. Empirical research indicates that attachment styles may be more stable in some contexts than in others.

In the Fraley et al. (2011b) studies people's attachment styles toward their parents and their romantic partners were assessed separately. When people were evaluating their attachments to their parents (mother and father, separately), the test-retest stability of attachment was over 0.80 (see the second row of Fig. 6.5). When people were evaluating their attachments to their romantic partners, however, the overall association was lower, closer to 0.60, on average (see the first row of Fig. 6.5). Although the patterns of correlations in both cases were consistent with a trait-state model, the overall level of stability was lower in romantic contexts than in parental contexts.

Why might there be more stability in parental than in romantic relationships? We believe Bowlby's ideas on canalization may be helpful for explaining these findings. The relationships adults have had with their parents have existed for decades. As a result, it seems likely that adults have settled into a robust pattern of interacting with their parents. Adults are no longer trying to gauge whether their parents are available, supportive, and responsive; they *know* whether their parents are relatively available and responsive. In addition to this, most adults are unlikely to interact with their parents on a daily basis. As a consequence, there are fewer opportunities for parental interactions that might lead to changes in the way in which one conceptualizes the parental relationship. These dynamics, of course, could change as people's parents' age and adult children become increasingly concerned about managing care for their aging parents.

In contrast, romantic relationships, by necessity, are often more nascent than the relationships people have with their parents. Not only are adults more likely to interact with their partners more than their parents, but also those relationships are still forming. Thus, as new events take place (eg, marriage, parenthood), there are multiple opportunities for those interaction patterns to shift around in subtle ways.

Although we believe the different patterns of stability observed in parental and romantic relationships are consistent with Bowlby's ideas concerning canalization, we should note that the research to date provides, at best, an indirect test of the key ideas. What would be ideal is an evaluation of trait-state models in the context of romantic relationships as they develop. The model implies that stability should be higher in long-term relationships than in short-term relationships. But evaluating this prediction in a truly compelling way would require examining the stability of attachment in young relationships as they develop across time.

DO ADULT ATTACHMENT STYLES CHANGE ACROSS THE LIFESPAN? NORMATIVE SHIFTS IN ATTACHMENT

Up to this point, our discussion of continuity and change has largely concerned the stability of individual differences: whether people who are relatively secure at one point in time are also likely to be relatively secure at another point in time. This is often referred to as “rank-order stability” in the personality literature because the primary concern is whether the relative ordering of people is the same across time. Another important form of stability, however, concerns mean-level or absolute stability. This is relevant to understanding whether, on average, people tend to increase (or decrease) in security across time. These two forms of stability are conceptually and mathematically independent of one another because people could preserve their rank ordering perfectly across two time points even if everyone became more secure, on average. And, similarly, even if the average levels of security were the same across two time points, if the people who were most secure at time 1 became the least secure at time 2 (and vice versa), mean-level stability could be perfect despite rank-order stability being zero.

One of the largest studies to examine mean-level stability and change in adult attachment was published by Chopik et al. (2013). They examined age differences in attachment style using a cross-sectional design in a sample of over 23,000 individuals ranging in age from 18 to 70. Attachment was assessed using the ECR-R, a self-report measure that focuses on romantic attachment, but not on specific romantic partners. Chopik and colleagues found that older individuals tended to have lower levels of attachment-related anxiety than younger individuals. One potential explanation for this finding is that, as people get older, they may have fewer reasons to be concerned with the availability and responsiveness of others. Chopik and colleagues also found that older individuals tended to be slightly more avoidant than younger individuals. A similar finding was reported by Magai et al. (2001). In a community sample of approximately 800 North Americans (average age of 74), they found that the average scores for avoidance were higher than they typically are in younger samples. Magai et al. (2001) attribute this finding to the long-term impact of economic hardship on families earlier in the century.

Hudson, Fraley, and Chopik (2015) conducted a similar analysis in a separate sample, but focused on attachment in specific relational contexts. That is, they assessed people’s general attachment orientation in addition to how people related specifically to their parents, their romantic partners, and their best friends. They found that people generally became less anxious with respect to attachment across time. That is, younger adults had higher anxiety scores than older adults. In contrast, there were few age differences in avoidance. For the most part, global avoidance tends to be relatively stable across age groups. The patterns of age-related differences varied across specific relational contexts, however. Younger people, for example, were generally more anxious in

romantic and friend relationships than older people. But the reverse was true in parental relationships. In parental relationships, younger people were less anxious than older people in their relationships with their parents. Why might this be the case? One possibility is that, as people's parents' age, people become less confident in the availability and responsiveness of their parents, potentially heightening the sense of anxiety people feel in their relationships with their parents.

In both peer (romantic and friendship relationships) and parental relationships, people seemed to become more avoidant across time. That is, older people were more avoidant toward their partners, friends, and parents than younger adults. The authors speculate that one reason for this shift is that role norms for adults typically emphasize a greater need for autonomy and independence as people make the transition from young to middle adulthood. It is also possible that the increase in avoidance in romantic relationships mirrors shifts in marital satisfaction that are commonly observed in long-term marriages. It is important to note that global avoidance, however, did not show systematic, replicable age differences across time. This suggests that, as a general rule, people do not become more avoidant across time, but the dynamics of specific relationships may create a press for greater degrees of independence with age. The obvious limitation of these studies is that they are based on cross-sectional data. To fully understand how attachment tends to change as a function of the life course, one needs to study people longitudinally as they develop.

SUMMARY AND FUTURE RESEARCH DIRECTIONS

Some of the fundamental questions in the study of attachment concern the stability of individual differences in attachment style: How stable are individual differences in attachment? What processes promote continuity and change? How does stability vary across relationship contexts and across different phases of the adult lifespan? We believe there are few broad conclusions that can be reached based on existing research. First, attachment styles appear to function in both trait-like and state-like ways. Although attachment styles tend to reflect variation in people's ongoing interpersonal experiences, underlying that variation are relatively stable dispositions—something that appears to undergird variation across time.

A second theme is that people's developmental trajectories appear to become increasingly canalized across time. Holding the time-interval constant, the test–retest stability in attachment observed early in romantic relationships is lower than the test–retest stability observed later. Moreover, people's representations of their parental relationships—relationships that are more established—are more stable than their representations of their romantic relationships. Taken together, these kinds of findings suggest that, in the early phases of a relationship (whether it be an infant–parent relationship or a fledgling romantic relationship), people construct working models of that relationship based on their

relational experiences. But, as the relationship progresses, those patterns of interaction begin to stabilize and working models begin to consolidate to some extent. The consequence is that people's sense of security or insecurity in a relationship is more resistant to change later than early in relationship development (Fraley & Brumbaugh, 2004).

A third theme is that people's working models, despite functioning as dispositional variables in adulthood, can and do change. On occasions in which people experience interpersonal losses, for example, they are more likely to feel insecure than they are on other occasions (Davila & Sargent, 2003). Certain developmental experiences, such as a history of family or personal psychopathology, have the potential to heighten the instability of working models, making people more likely to change, but not always nudging them in a specific direction consistently (Davila et al., 1997). Moreover, above and beyond changes in the rank ordering of individual differences, cross-sectional research suggests that people tend to become less anxious across time in their romantic relationships (see also chapter: What Are the Effects of Context on Attachment?).

Despite the progress that has been made in understanding stability and change in attachment, we still have a lot of work to do to more fully understand the dynamics of stability and change. One of the gaps in our current knowledge is that we do not have a strong handle on how specific experiences impact attachment because we tend to know very little about what people are like before certain events or transitions take place in their lives. Take loss as an example. There is a large literature on how losing a loved one can disrupt psychological functioning (Parkes & Weiss, 1983), and many scholars have been interested in the question of whether attachment style, for example, predisposes people to experience chronic or disordered forms of grief (Fraley & Bonanno, 2004). But much of the research that examines the association between attachment and adaptation to loss has been forced to assess bereaved people's attachment styles *after* the loss has taken place. As a consequence, it is difficult to know whether people are responding in ways—adaptive or maladaptive—that are predictable from their preloss attachment orientation and how and if that orientation is affected by the loss itself.

One solution to this kind of problem is to assess attachment styles across multiple occasions before major life events take place. This would enable researchers to characterize the person's prototypical trajectory of attachment (ie, the extent to which the person's pattern is state-like and trait-like, whether he or she is increasing in anxiety over time and the rate at which he or she does so, and the amount of variability the person exhibits across time) before specific events take place and to examine the ways in which specific events alter or disrupt that trajectory.

A second gap in the existing literature involves understanding whether certain experiences have short- or long-term consequences for attachment orientation. At first glance, this would seem like a relatively simple problem to solve. All a researcher would need to do, presumably, is examine the consequences

of a specific event (eg, a breakup, the death of a loved one) in a long-term follow-up study. But as Fraley, Roisman, and Haltigan (2013) observed, these kinds of designs do not enable one to determine whether a specific event had enduring or transient consequences for the outcome in question. Imagine, for example, that a research team finds that the association between psychotherapy and attachment security assessed 6 months later is 0.30. Most researchers would conclude on the basis of those data that therapy has long-term benefits for security. But what if the team were to continue their assessments and learn that the association between treatment and security is 0.10 after 8 months, and 0.00 after 12 months? That would lead to a dramatically different conclusion about the efficacy of therapy than if we observed, instead, that the association between treatment and security was 0.30 after 6, 8, and 12 months. What are needed in future research are multiple measurements of the construct across time so one can detect whether the changes are persistent and sustained or whether they are getting smaller as the time between the event and the outcome increases.