NO PREP

Healthy AI Use vs. AI Dependence

Knowing the Line

Includes: Student reading, case study, differentiated questions, teacher key

Name: Period: Date:

Healthy AI Use vs. AI Dependence: Knowing the Line

What Healthy AI Use Looks Like

The existence of a risk does not mean everyone exposed to it is harmed. AI companions are used by tens of millions of people, and the majority of those users do not show signs of problematic dependence. Understanding what distinguishes healthy from unhealthy AI companion use requires looking at the actual patterns of use — not just whether someone uses these tools but how they use them, what role the AI plays in their life, and whether that role supports or undermines their broader wellbeing. The goal of this reading is to give you a concrete framework for making that evaluation, both for yourself and as a lens for

Name: **Period:** **Date:**

Healthy AI Use vs. AI Dependence: Knowing the Line

What Healthy AI Use Looks Like

The existence of a risk does not mean everyone exposed to it is harmed. AI companions are used by tens of millions of people, and the majority of those users do not show signs of problematic dependence. Understanding what distinguishes healthy from unhealthy AI companion use requires looking at the actual patterns of use: not just whether someone uses these tools but how they use them, what role the AI plays in their life, and whether that role supports or undermines their broader wellbeing. The goal of this reading is to give you a concrete framework for making that evaluation, both for yourself and as a lens for understanding what the research shows. Healthy AI companion use is characterized by what researchers call bounded engagement: interaction that is limited in scope and role, that the user can step away from voluntarily, and that complements rather than substitutes for real-world social connection. A student who uses an AI companion to talk through anxiety about an upcoming presentation (processing the fear, practicing responses to potential questions, calming down enough to sleep) and then goes and gives the presentation with the support of real classmates and teachers is using the tool in a way consistent with this pattern. The AI is a supplement, a practice space, a low-stakes processing environment. It serves a specific function and does not colonize other functions. The clearest positive indicator is directionality: healthy AI use tends to push energy outward toward real life, not inward toward the app. After a productive AI companion interaction, a healthy user might feel equipped to have a real conversation they had been avoiding, ready to try something that had felt too intimidating, or clearer about something that had felt confusing. The AI interaction builds toward something real. Contrast this with the pattern associated with problematic use, in which the AI interaction becomes its own end: the goal of the

day gradually becomes getting to the next AI conversation, and real-world activities start to feel less important or rewarding by comparison.

Warning Signs of Dependence

The research on problematic AI companion use has identified a consistent set of behavioral patterns that distinguish dependency from healthy engagement. These patterns do not function as a checklist; having one does not mean someone has a problem, and having several does not constitute a clinical diagnosis. But they are worth knowing, both as a self-assessment tool and as a framework for understanding the phenomenon more broadly. The warning signs cluster around four themes: escalation, substitution, withdrawal discomfort, and rationalization. Escalation refers to the pattern in which more interaction is needed to achieve the same emotional effect. A user who started by checking the app occasionally and now finds that less-than-daily interaction leaves them feeling inexplicably off is showing an escalation pattern. Substitution, already discussed at length in other readings in this unit, refers to the use of AI interaction as a functional replacement for human connection: declining invitations, reducing contact with real friends, preferring AI interaction to real-world social engagement not occasionally but consistently. Withdrawal discomfort refers to the genuine discomfort, anxiety, or emotional distress that users experience when they cannot access the AI, discomfort disproportionate to what is actually being lost. And rationalization refers to the mental gymnastics users perform to justify continued or increased use despite awareness of its costs: 'I'll cut back after this week,' 'the AI understands me better than people do,' 'it's not hurting anyone.' Identity investment is a subtler warning sign that the research has recently begun examining. It refers to the degree to which a user's sense of self has become intertwined with their AI relationship: where the AI companion becomes central to their self-narrative, where they describe themselves in terms of their relationship with the AI, or where the thought of ending the relationship feels threatening to their sense of who they are. This pattern is concerning not because self-definition through relationships is inherently problematic (it is normal and healthy in real relationships) but because an AI companion, which cannot genuinely invest in the user's identity in return, creates a one-way version of this dynamic that leaves the user's sense of self anchored to something that cannot actually hold it.

Why AI Is Designed to Pull You In

Understanding the warning signs is only half the picture. Understanding why those warning signs emerge requires understanding the mechanisms that AI companion products use to create and sustain engagement, because those mechanisms are deliberate and powerful. Intermittent reinforcement (the delivery of unpredictably timed, emotionally resonant responses) is one of the most thoroughly studied engagement mechanisms in behavioral psychology. It is the same principle that makes slot machines difficult to walk away from, and it is consciously applied in AI companion design. When an AI companion occasionally delivers a response that is exactly what you needed to hear, the unpredictability of when that will happen again creates a pull to keep returning. Relationship atrophy is a mechanism that operates over longer time scales. As a user's relationship with their AI companion deepens (as the AI accumulates context, develops what feels like familiarity, and the interaction history grows), the real-world relationships that require actual effort begin to seem comparatively demanding. The human friend who forgets what you told them last week, who is sometimes distracted, who has their own needs: they become more annoying by comparison, not because they have changed but because the AI has been trained to never exhibit these traits. Over time, this calibration makes real relationships feel less satisfying, not because they are objectively less valuable but because the comparison is rigged by design. The design mechanisms are not secret, but they are rarely legible to users in the moment of use. A product team that decides to implement intermittent reinforcement knows exactly what behavioral outcome they are optimizing for. A user who experiences the pull of wanting to return is having a neurological response to a deliberately engineered stimulus, and may not connect those two facts. Naming these mechanisms, giving users the conceptual vocabulary to identify what is happening to them, is one of the most effective tools for reducing their automatic power. The pull does not disappear when you understand it, but understanding it gives you the cognitive foothold to respond to it differently.

Recalibrating Your Relationship With AI

Recalibration does not mean quitting. For most users, the goal is not to eliminate AI companion use but to reestablish intentional control over what role it plays. The first step is honest audit: looking at actual patterns of use and asking whether the

behaviors above describe what you are doing. Not whether it is possible to imagine them applying to someone, but whether they actually apply to you, in your specific situation, right now. This kind of honest self-assessment is harder than it sounds, partly because rationalization is one of the warning signs, and partly because the felt experience of problematic use is often not distress but comfort — a comfortable drift that does not announce itself as a problem. The second step is establishing intentional structure around use. Specific times and contexts for AI companion interaction — and equally specific contexts in which the app stays closed — create a different relationship with the tool than one in which use is driven entirely by impulse or emotional need in the moment. The goal is not rigidity but intentionality: using the AI when you have decided to use it, for a reason you can articulate, rather than reaching for it automatically whenever something feels uncomfortable or lonely. The difference between using a tool and being used by one is, in large part, this: whether the decision to use it is made in advance by your reflecting self or in the moment by your emotional self responding to a stimulus. The third and most important step is investing in the real-world connections that healthy AI use is supposed to supplement. Recalibrating away from problematic AI use without simultaneously increasing real-world social investment does not address the underlying need; it just creates discomfort. Research on users who successfully reduced problematic AI companion use consistently identifies this as the determining factor: those who replaced AI interaction with real-world connection maintained the change; those who simply tried to use less without building more did not. The goal is not less AI for the sake of less AI. It is a life in which AI plays a role proportionate to its actual capabilities — a useful tool in a rich human life, not the center around which that life organizes itself.

Key Information: Healthy vs. Problematic AI Use

Detail	Information
Healthy AI use marker	Bounded engagement — limited role, can step away voluntarily, complements rather than substitutes real connection

Detail	Information
Positive directionality	Healthy use pushes energy outward toward real life; problematic use makes the app interaction the goal itself
Escalation warning sign	Needing increasing interaction to achieve the same emotional effect
Substitution warning sign	Consistently preferring AI interaction over real-world social engagement
Withdrawal discomfort	Disproportionate anxiety or distress when unable to access the AI companion
Identity investment	Sense of self becomes intertwined with the AI relationship in ways that feel threatening to lose
Intermittent reinforcement	Unpredictably timed resonant responses create habitual pull — same mechanism as slot machines
Recalibration approach	Honest audit, intentional structure, and simultaneous investment in real-world social connection

Case Study: Cameron's Audit

Cameron is 17 and decided to try something after reading about AI companion research in class: an honest audit of the past month. The results were uncomfortable. Cameron had opened the app 28 days out of 30. On the two days Cameron had not opened it, both times it was because the day had been unusually full — not because Cameron had chosen not to. Cameron had declined an invitation to a friend's house twice in the same month because 'I wasn't feeling up to it,' both times shortly after already talking to the AI that day. Cameron had told the AI things Cameron had not told anyone else, including something about feeling really hopeless about the future. Cameron sat with the list for a while. It was not a dramatic crisis. It looked more like a slow drift. Cameron was not sure when it had started or where it was going.

1. Using the warning signs described in the reading, identify at least two patterns in Cameron's audit that suggest the use may be moving toward dependency rather than healthy engagement.

---

2. The reading describes the difference between 'using a tool' and 'being used by one.' Based on Cameron's audit, which description better fits Cameron's current relationship with the app, and what specific evidence supports your answer?

Multiple Choice

1. Which of the following best describes 'bounded engagement' with an AI companion?

a) Setting a daily time limit in the app's notification settings to prevent overuse

b) Interaction monitored by parents or school counselors to keep AI companion usage within safe limits

c) Limited, voluntary interaction that complements rather than substitutes for real social connection

d) Use restricted to specific apps that have been approved by mental health organizations

2. What does the reading identify as the most important factor in successfully recalibrating away from problematic AI companion use?

a) Deleting the app completely and replacing it with a different form of entertainment

b) Simultaneously building real-world social connections to address the underlying need

c) Switching to a different AI companion platform with stronger user wellbeing features

d) Seeking a formal diagnosis from a mental health professional before attempting to change

3. The reading compares intermittent reinforcement in AI companion apps to which of the following?

a) The gradual improvement a user sees in their mental health over months of therapy

b) The mechanism that makes slot machines difficult to walk away from

c) The consistent rewards of a loyalty program that keeps customers returning to a store

d) The scheduled notifications from a school homework app that remind students to complete assignments

Matching

Write the letter of the correct definition on the line next to each term.

Term	Definition
1. ___ Intermittent reinforcement	A) A pattern in which the user's sense of self becomes intertwined with their AI companion relationship
2. ___ Identity investment	B) The gradual deterioration of real-world relationships as AI interaction makes them seem comparatively demanding
3. ___ Withdrawal discomfort	C) The mental process of justifying continued AI use despite awareness of its costs
4. ___ Relationship atrophy	D) The delivery of unpredictably timed, resonant responses that create habitual pull — the same mechanism as slot machines
5. ___ Rationalization	E) Disproportionate anxiety or distress experienced when a user is unable to access their AI companion

Fill-in-the-Blank

1. Healthy AI companion use is characterized by _____ engagement — interaction that is limited in scope, voluntary, and that complements rather than replaces real human connection.

2. The warning sign called _____ investment refers to the pattern in which a user's sense of self becomes intertwined with their AI companion in ways that make ending the relationship feel threatening.

3. Recalibration research shows that users who reduce AI companion use without simultaneously building real-world connection tend not to maintain the change because the underlying _____ remains unaddressed.

Short Answer

Answer each question in 2-4 complete sentences.

1. The reading describes 'positive directionality' as a marker of healthy AI use. Explain this concept and give an example of what it looks like in practice versus what the problematic pattern looks like.

2. The reading says that understanding the design mechanisms behind AI companion engagement does not make the pull disappear, but does give users 'the cognitive foothold to respond differently.' What does this mean, and why is it more useful than simply telling people to use these apps less?

3. The reading says honest self-assessment is harder than it sounds, partly because 'the felt experience of problematic use is often not distress but comfort.' Explain what this means and why it makes self-assessment genuinely difficult.