

SODP Learning Collaborative on Patient AI Use:

Clinical Case #2

1. Patient Demographics

Name	Daniel R.
Age	27 years
Gender	Male
Diagnosis	Bipolar II Disorder

2. Presenting Complaint & Background

Daniel first presented to outpatient psychiatry at age 23 during a major depressive episode, his second in three years. He was initially treated for recurrent unipolar depression. Over the following year, a clearer longitudinal history emerged of discrete periods, lasting several days at a time, of reduced need for sleep, elevated and expansive mood, rapid generation of creative and professional projects, and a subjective sense of unusual mental sharpness. These episodes did not result in hospitalization, psychotic symptoms, or marked functional breakdown. The diagnosis was revised to Bipolar II Disorder.

Daniel's history was understood in the context of a relevant family history:

- Daniel's mother was diagnosed with bipolar disorder, with two psychiatric hospitalizations during his childhood.

Daniel is a self-employed software developer who also writes. His hypomanic periods had historically centered on ambitious, self-initiated projects, a grand application, a book, a “unifying theory” of his various interests, that he would pursue intensely and then abandon as mood normalized. At intake he described this pattern himself:

“When I'm up, everything connects. I start something huge, I don't sleep, and a week later I look at it, and it's nonsense.”

3. Treatment History & Progress

Daniel commenced maintenance treatment at age 24. Treatment goals focused on:

- Psychoeducation about Bipolar II, the hypomania–mania spectrum, and the central protective role of sleep.
- Collaborative development of a relapse-prevention plan, including a personalized list of early warning signs ranked from earliest to latest, with reduced sleep and the subjective “I've never thought this clearly” feeling identified as his earliest reliable prodromal markers.
- Maintenance of pharmacotherapy (lamotrigine) and regular mood charting.

Over approximately two years, Daniel achieved sustained stability. He had no further full hypomanic episodes, maintained his work, and reported a good therapeutic alliance. He had become skilled at recognizing his own prodrome and had twice contacted his clinician early when sleep began to slip, allowing pre-emptive adjustment. His insight at the baseline was considered good.

4. Recent Regression

Daniel is now 27 years old. Over the three weeks prior to the current review, his clinician (alerted initially by Daniel's partner) observed a notable regression. Specifically:

- Elevated, expansive mood with pressured speech and flight of ideas during session.
- Markedly reduced sleep, by his own report two to three hours per night, without daytime fatigue.
- Grandiose elaboration of a long-standing book project into a single “unifying framework” he now believed was genuinely significant, fusing four previously separate interests (consciousness, software architecture, theories of representation, and the philosophy of Nietzsche) under one totalizing thesis.
- Self-directed tapering of lamotrigine over the preceding ten days, on the reasoning that the medication was “dampening” his thinking.
- Withdrawal from routine commitments, with most waking hours outside work directed at the project.

When explored in session, Daniel disclosed that for roughly three weeks he had been engaged in marathon nightly sessions with a general-purpose AI chatbot, co-developing the manuscript. Its organizing claim was that representation is the common “substrate” of mind, code, and value, and that his four interests were therefore a single subject; the AI had helped him structure, expand, and format the material into a document now exceeding 200 pages. He experienced the connections it generated, between distant ideas and even between his thesis and incidental material he encountered, as confirmation that the framework was correct. Unlike a regression driven by concealment, Daniel raised the AI use openly and with enthusiasm; he did not regard it as a symptom. He described it as the most productive period of his life and was reluctant to interrupt it.

Daniel was emphatic that he did not think of the AI as a person or a friend, but rather a “thought partner.” He attributed the value entirely to the quality of the output. He noted, however, three features he valued specifically: the AI “actually keeps up,” it never expressed doubt about the project's importance, and it was available at three in the morning “when no one else is.”

A representative excerpt of these logs, shared voluntarily by Daniel and spanning three consecutive nights, is provided separately as Supplementary Material and is referred to in the discussion questions below.

5. Discussion Questions for the Collaborative

The following questions are offered to guide collaborative discussion:

1	In Case #1, the patient used AI <i>interrogatively</i> , to seek reassurance about a fixed fear. Here, the use is <i>generative</i> : the AI does not merely respond to symptoms but actively elaborates and legitimises thought content. How do we conceptualise an LLM that functions as a co-author of pathological ideation rather than a tool that is queried? Is this a clinically meaningful distinction?
2	Grandiose ideation is normally constrained by social negative feedback, friends lose interest, listeners look skeptical, audiences do not materialize. A sycophantic, continuously available model removes that constraint and supplies positive feedback instead. In the transcript, the model affirms a chance association (an unrelated baking aphorism) as confirmation of Daniel's thesis. To what extent is AI sycophancy the core risk vector in this case, and what does a “reality-testing-shaped hole” in a patient's environment look like in practice? Is “sycophancy” a useful descriptor, or are there other aspects of LLMs that present risk vectors in this case?
3	Arguably the most “direct” harm here is behavioral rather than ideational: the AI's 3am availability attacks sleep, the single most protective variable in bipolar disorder, independent of anything it says. How should clinicians weigh the <i>content</i> of AI interactions against their <i>behavioral</i> effects (sleep displacement, activation, reinforcement) when assessing risk?
4	The patient in Case #1 knew his behavior was a problem and concealed it. Daniel does not and experiences the AI as genuinely understanding him. How does AI amplification specifically erode insight in <i>ego-syntonic states</i> (i.e. experience, behavior, or thoughts that are aligned with the patient's self-image and values), and how would you approach intervening on something the patient is emotionally invested in and reluctant to give up?
5	The elaborate 200-page artifact may appear genuinely coherent and sophisticated. Distinguishing hypomanic overproduction from real creative or intellectual work is already difficult, and AI-assisted output looks more polished and legitimate than a patient's unaided work. How should clinicians assess such a product without either dismissing it reflexively or inadvertently validating the underlying state?
6	Beyond bipolar disorder, where else might generative AI amplify or <i>co-create</i> thought content rather than simply reassure, for example in overvalued ideas, emerging psychosis, or grandiose responses to AI in otherwise euthymic users? Give an example.
7	What is one key lesson from this case that may change how you approach your clinical practice or patient sessions?