Advanced Risk Formulation and Safety Planning Practice

THE TALKING THERAPIES UK PRACTITIONER SERIES



Advanced Risk Formulation and Safety Planning Practice

Complex and chronic presentations, supervision and organisational risk policy

Advanced level · 16.5 hours of continuing professional development

TALKING THERAPIES UK

First edition · 2026

Published by Talking Therapies UK, talkingtherapies.co.uk.

Copyright © 2026 Talking Therapies UK. All rights reserved.

The moral right of the author has been asserted. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in a review, and except that a purchaser of the accompanying course may print a single copy for their own study.

First edition, 2026.

This manual accompanies the Advanced Risk Formulation and Safety Planning Practice course in the Talking Therapies UK Digital Classroom, from which its content is drawn and against which it is kept current.

Revision 4. This printing was set on 22 November 2026 from the course as it stood that day.

A note on the clinical material. Every case in this book is a fictional composite written for teaching. No real person appears in it, however disguised. Nothing here constitutes clinical supervision, and nothing in it replaces the judgement of a qualified practitioner working within their competence, their professional body's code and their local governance. Where guidance is cited it was current at the date of this edition; readers must check the standing of any guidance before relying on it.

Set in Poppins and Roboto, the Talking Therapies UK typefaces. Figures drawn by the Talking Therapies UK teaching-diagram engine.

Contents

How to use this manual 1

Part 1 Formulating Chronic and Recurrent Risk Presentations 3

1. Working with chronic or recurrent presentations 3
2. Working across settings and with family or support networks 10
3. Transitions, Endings and Handovers in Complex Chronic Risk Work 17
4. Distinguishing chronicity from acuity: formulating risk when crisis becomes a recurring pattern rather than a single episode 24
5. A functional analysis of recurrent self-harm: understanding what the pattern does for the person, not just what it risks 31

Part 2 Sustaining Collaborative Practice in Long-Term Risk Work 38

6. Avoiding alarm fatigue and complacency over a long course of chronic risk work 38

How to use this manual

This advanced course addresses the presentations and organisational questions that go beyond a single collaborative conversation: chronic or recurrent patterns of self-harm, coordination across multiple professionals and settings, supervision of this work, and organisational risk policy. You will study how formulation and safety planning need to adapt for someone with a longstanding pattern of crises rather than a single acute episode, guarding against the therapeutic fatigue that can creep in when a pattern feels familiar, and updating a formulation as circumstances genuinely change rather than reusing an old one unreflectively. The course then teaches NICE NG225's recognition that self-harm presents across many settings beyond mental health services, including primary care, education and criminal justice, and the role of family involvement where the person wants it, before turning to what supervision of this work should specifically check for and how to build organisational policy that reflects NG225's current position rather than legacy risk-stratification practice. The course is intended for senior practitioners and those responsible for supervision or policy in this area, and sits within the practice's existing professional practice and supervision pathway.

This manual is the written form of a course taught in the Talking Therapies UK Digital Classroom. It is set at advanced level and carries 16.5 hours of continuing professional development for practitioners who complete the accompanying course and its assessment.

It is written to be read in order. Each part opens with what it is for; each chapter teaches one idea to a working depth and closes with the points worth keeping and a practice task that converts reading into something you do. The tasks are the part most often skipped and the part that most reliably changes practice, and several of them are designed to be taken to supervision rather than completed alone.

The figures are teaching diagrams rather than decoration. Each is the conventional form of a model you will meet in the literature and in supervision, drawn in one visual language across the whole series so that the shape of an idea is recognisable wherever it reappears.

Where a chapter draws on published work, the reference is named in the text and gathered in the further reading at the back. Terms of art are defined in the glossary, which is worth reading straight through once before you begin.

What this manual will leave you able to do

- Recognise and guard against therapeutic fatigue in chronic or recurrent risk presentations
- Update a risk formulation as circumstances change rather than treating it as fixed
- Involve family or support networks in risk work only where the person themselves wants it
- Supervise risk formulation and safety planning work, checking specifically for drift back into risk-scoring
- Attend to practitioner wellbeing within supervision of sustained risk work
- Review organisational risk policy against NICE NG225 and maintain consistent team training

PART 1

Formulating Chronic and Recurrent Risk Presentations

This module examines how to assess, understand and formulate risk when crisis becomes a recurring pattern, including functional analysis of self-harm and working across settings and networks.

CHAPTER 1

Working with chronic or recurrent presentations

Working with chronic or recurrent presentations represents one of the most demanding yet critical facets of clinical practice within the UK Talking Therapies sector. Unlike acute presentations, where the crisis is a singular, often precipitous event, chronic risk involves a long-term pattern of distress, self-harm, or suicidal ideation that has persisted over months or years. This continuity fundamentally alters the therapeutic dynamic, shifting the focus from acute crisis intervention to the management of a persistent, fluctuating threat. For the practitioner, this requires a sophisticated adaptation of risk formulation and safety planning, moving beyond the linear models often taught for initial assessments to embrace a more fluid, iterative process that acknowledges the client's history while remaining vigilant to the present moment. The central challenge in this work is the phenomenon of therapeutic fatigue, a state of exhaustion that can erode the therapeutic alliance if not actively managed by both clinician and client. This fatigue is not merely emotional burnout for the therapist; it is a psychological state for the client who may feel trapped in a loop where their efforts to manage their

mental health yield diminishing returns, leading to a profound sense of stagnation and hopelessness.

The specific risk of therapeutic fatigue with recurrent presentations is that it creates a barrier to engagement; the client may arrive at sessions feeling that they are merely repeating a script they have already written, and the therapist may unconsciously feel that they are rehashing old ground. This dynamic is dangerous because it can lead to a subtle erosion of the therapeutic alliance, where the client feels seen but not heard, and the therapist feels competent but not challenged. The antidote to this fatigue is a rigorous commitment to the principles outlined in the National Institute for Health and Care Excellence (NICE) guidelines, specifically NG225, which emphasises a compassionate, non-judgmental approach that is consistent, yet this consistency must be paired with a disciplined intellectual vigilance that refuses to accept the status quo. Familiarity should never translate into complacency. The assumption that “they have been through this before and will be fine” is a dangerous cognitive distortion that can lead to underestimation of lethality. When a client presents with a crisis that mirrors a previous one, the natural human tendency is to rely on established patterns of thinking and responding. This is where the danger lies.

In the context of chronic risk, the “same” crisis can be a manifestation of a deeper, more complex set of vulnerabilities that have evolved over time. The evidence base regarding repeated self-harm suggests that while clients may develop better coping strategies over time, their underlying vulnerabilities often remain, sometimes evolving into more complex or lethal presentations. Models such as Cognitive Behavioural Therapy (CBT) and Schema Therapy provide frameworks for understanding these cycles. In CBT, the focus might be on the maintenance cycle of automatic thoughts and behaviours, but in chronic cases, we must look at the “maintenance of the problem” rather than just the “problem itself.” The Cognitive Triad of depression—negative views of the self, the world, and the future—remains a potent predictor of chronic suicidality, but its expression changes. A client who previously self-harmed to regulate overwhelming emotion may now use more lethal means due to a perceived increase in hopelessness. This requires the practitioner to constantly re-evaluate the formulation. We must be precise with our terminology. Tools such as the Social GGRRAAACCEEESSS or the PHQ-9 are invaluable, but they are static snapshots. If a client completes a PHQ-9 every session and the score remains high, the tool itself is not providing new

information; it is merely confirming a chronic state.

The danger is that the therapist stops looking for the change in the variables that drive the score. The Social GGRRAAACCEEESSS, while excellent for mapping social support, can become a stale document if not updated to reflect new relationships, new conflicts, or the withdrawal of support. We must treat these tools as living documents that require active, ongoing maintenance. The discipline of updating a formulation is paramount. A formulation is not a diagnosis; it is a hypothesis about the function of behaviour. When a client has a chronic presentation, the hypothesis must be tested and refined regularly. This involves looking for “triggers” that have shifted. Perhaps a client previously self-harmed during holidays but now self-harms on weekends when alone. This shift indicates a change in their social environment or internal regulation capacity. The formulation must capture this nuance. If we default to an old formulation simply because it is easier, we risk “formulation drift,” where the clinical picture becomes detached from reality. We must ask, “What has changed?” rather than “What is the same?” Maintaining compassionate engagement across repeated contacts is the antidote to fatigue. When a client returns for their fifth crisis in a year, they may feel like a burden.

The therapist must convey that the relationship is not transactional but is built on a deep, enduring commitment. This requires active listening and a willingness to sit with the discomfort of repetition without becoming bored or dismissive. The therapist must model resilience and hope. If the therapist shows fatigue, the client will feel that their distress is not worth the effort, potentially leading to an increase in risk. To illustrate these principles, consider the case of Sarah, a client with a ten-year history of recurrent suicidal ideation and self-harm. In her first year of therapy, Sarah’s formulation was heavily focused on childhood trauma and low self-esteem. Her safety plan was standard. Three years later, Sarah lost her job. Her formulation needed updating. The trigger was no longer just internal; it was external. Her social support network had shrunk. Her safety plan needed to include specific strategies for unemployment-related anxiety. Five years later, Sarah met a new partner. Her formulation shifted again. The risk factors changed from isolation to relationship conflict. The lethality of her self-harm increased because she was using more dangerous methods.

The therapist had to update the formulation to reflect that her “self-harm

was a way to manage relationship rejection," not just "general distress." The safety plan had to be revised to include conflict resolution skills and specific contact numbers for crisis support. Had the therapist relied on the original formulation, they might have missed the specific link between her relationship status and her increased risk, leading to a potential tragedy. This case demonstrates that in chronic work, the formulation is a dynamic map, not a static destination. It requires constant navigation, regular recalibration, and an unwavering commitment to the belief that even in the most repetitive of patterns, there is always a new variable waiting to be discovered.

Returning to Sarah's case, we see that as her relationship progressed, the therapeutic team recognised the need for a more specialized approach to relationship-based risk factors. The therapist initiated a collaborative review session where Sarah was invited to map the timeline of her relationship alongside her self-harm episodes. This visual exercise revealed a pattern: her self-harm occurred most frequently during periods of perceived rejection or abandonment, even when these perceptions were objectively unfounded. The formulation was updated to incorporate Schema Therapy concepts, identifying schemas such as abandonment and emotional deprivation as core drivers of her risk. The safety plan was expanded to include specific strategies for challenging these schemas during moments of relational distress, such as thought records and behavioural experiments to test her perceptions of rejection. This collaborative approach empowered Sarah to become an active participant in her risk management, transforming her from a passive recipient of care to an informed collaborator in her own safety.

A common pitfall in chronic risk work is what we term 'therapeutic drift'—the gradual erosion of clinical precision as cases become longer-term. This manifests when practitioners rely on established patterns of response without critically evaluating their ongoing relevance. For example, a therapist might continue to focus exclusively on Sarah's childhood trauma even when current relationship conflicts have become the primary risk factor. This drift often occurs subtly, with clinicians maintaining a facade of competence while their clinical reasoning becomes increasingly disconnected from the client's present reality. The antidote to therapeutic drift is disciplined clinical curiosity—actively seeking disconfirming evidence to one's hypotheses and maintaining a stance of 'not knowing' that invites new information. In Sarah's case, this meant the therapist regularly asked, 'What am I missing?' and 'What new understanding might emerge if I view this through a different

lens?'

Supervision plays a critical role in preventing therapeutic drift in chronic cases. When presenting Sarah's case in supervision, the therapist was encouraged to structure their reflections around three key questions: 'What has changed since our last supervision?', 'What am I avoiding discussing with Sarah?', and 'What aspects of the formulation feel stale or untested?'. This structured approach helped identify that the therapist had become overly focused on Sarah's self-harm behaviours while neglecting to explore the meaning she attributed to her partner's actions. The supervisor suggested exploring Sarah's attachment patterns more deeply, which revealed that her perception of rejection was often triggered by normal relationship behaviours that activated early abandonment experiences. This insight led to a more targeted intervention focusing on emotional regulation in relationship contexts rather than solely on crisis management.

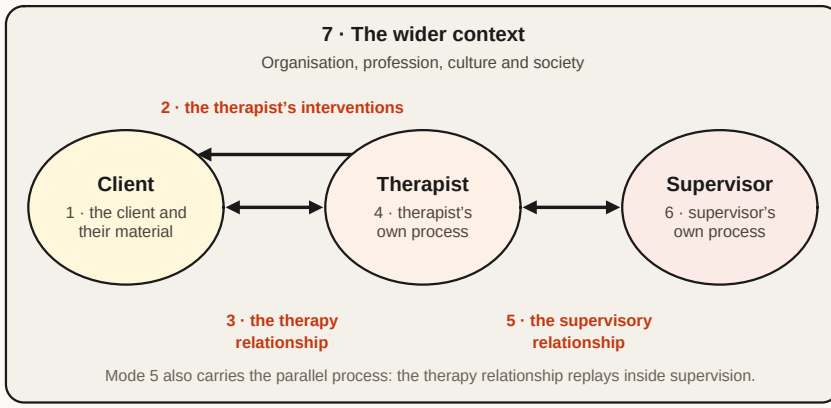
Regular formulation reviews should be scheduled systematically, not just when crises occur. For chronic cases like Sarah's, we recommend quarterly reviews that involve both therapist and client in a collaborative process. These reviews begin with the client's perspective: 'How has your understanding of your difficulties changed since we last met?' This question often reveals shifts in the client's self-awareness that may not be evident in session-by-session notes. The therapist then presents their evolving formulation, highlighting both continuities and changes. Crucially, these reviews should include a 'risk calibration' exercise—reassessing the relative importance of different risk factors based on recent data. In Sarah's case, this calibration revealed that while relationship conflict remained a significant trigger, her employment status had become a more potent risk factor following a recent job loss, necessitating a revised safety plan.

Maintaining hope in chronic cases presents a unique challenge. When a client has experienced multiple cycles of improvement and relapse, they may develop a 'hope fatigue' that mirrors therapeutic fatigue. The therapist must model a nuanced understanding of progress that acknowledges both setbacks and gains. For Sarah, this involved creating a 'progress timeline' that documented not just reductions in self-harm, but also improvements in her ability to identify triggers, use coping strategies, and seek help earlier in crisis cycles. This timeline became a powerful counter-narrative to her belief that 'nothing ever changes.' The therapist also normalised the process of

change, explaining that recovery from chronic conditions often follows a ‘two steps forward, one step back’ pattern, with each cycle building resilience even when setbacks occur.

Finally, we must address the ethical dimension of chronic risk work. Practitioners often experience moral distress when working with clients who remain at risk despite their best efforts. This distress can lead to either overprotective interventions that undermine the client’s autonomy or withdrawal of engagement that risks safety. The key is to maintain a balance between responsibility and respect for the client’s agency. In Sarah’s case, this meant negotiating boundaries around safety behaviours—for example, agreeing that while she would not completely eliminate self-harm, she would commit to using less harmful methods and maintaining communication with her therapist during high-risk periods. This collaborative approach honoured both the client’s autonomy and the therapist’s responsibility to prevent serious harm.

As we conclude this exploration of chronic risk formulation, it is essential to remember that the most effective approach is not one of rigid adherence to protocols, but of flexible, compassionate clinical reasoning. The case of Sarah illustrates how chronic risk work requires us to be both scientists—rigorously testing and updating our formulations—and artists—maintaining hope and connection in the face of repetition. The true measure of success in chronic risk cases is not simply the absence of crisis, but the client’s developing capacity to understand their own risk patterns and participate meaningfully in their safety. This is the transformative potential of risk formulation and safety planning in chronic presentations—not just managing risk, but empowering clients to become their own risk managers.



After Hawkins and Shohet.

FIGURE 1.1 The seven-eyed model of supervision

Hawkins and Shohet's seven modes give supervision its lenses: the client and their material, the therapist's interventions, the therapy relationship, the therapist's own process, the supervisory relationship itself (including parallel process), the supervisor's process, and the wider organisational and social context around all of it.

Points worth keeping

- The specific risk of therapeutic fatigue with recurrent presentations
- Why familiarity should never translate into complacency
- Updating formulation as circumstances genuinely change
- Maintaining compassionate engagement across repeated contacts

PRACTICE TASK

Describe, in two sentences, how you would check whether a long-held formulation for a recurrent presentation still fits the person's current circumstances.

CHAPTER 2

Working across settings and with family or support networks

Self-harm is not a phenomenon that presents neatly within the boundaries of specialist mental health services, and any advanced practitioner who works only from within those boundaries will encounter a distorted picture of the population they are trying to serve. People who self-harm are encountered by school counsellors and pastoral leads, by general practitioners and practice nurses, by emergency department staff, by youth workers, by probation officers, by housing support workers, and by many others who would not describe themselves primarily as mental health clinicians. NICE, in its guideline NG225, is explicit on this point: the guidance is addressed to the full range of settings in which self-harm presents, and it places a clear expectation on practitioners in all those settings to respond with the same quality of formulation-informed thinking that would be expected in a specialist service. For the advanced practitioner, this is not a peripheral consideration. It is central to understanding how risk formulation and safety planning actually function in the real world, where the person in front of you may already be known to several other professionals, each of whom holds a partial picture.

The clinical and ethical implications of this multi-setting reality are considerable. When a young person is simultaneously supported by a school counsellor, a CAMHS worker, a GP, and a youth offending team worker, each of those professionals is likely to be conducting some form of risk assessment, however informal. Each will be receiving disclosures, observing behaviour, and forming impressions. If those impressions are never shared, synthesised, or tested against one another, the result is not four partial pictures that together add up to a whole. The result is four separate and potentially contradictory narratives, each of which may be driving different responses. The school counsellor may be working on the assumption that the self-harm is primarily a response to peer difficulties, while the CAMHS worker is formulating around early attachment disruption, while the GP is focused on a recent medication change, and the youth offending team worker is managing risk of reoffending as their primary concern. None of

these perspectives is necessarily wrong, but without coordination they can pull the person in different directions, create confusion about who holds responsibility for what, and in the worst cases leave dangerous gaps through which serious harm can fall.

The advanced practitioner's role in this landscape is not simply to be one more voice in the crowd. It is to hold, or to help create, a coherent formulation that can travel across settings and remain recognisable as it does so. This requires a particular kind of clinical discipline. A formulation that is expressed in highly technical language, or that is so lengthy and nuanced that it cannot be communicated in a brief handover, will not survive contact with a busy emergency department or a primary care consultation of ten minutes. The formulation needs to be distilled into a form that preserves its essential meaning while being accessible to professionals with different training backgrounds and different time constraints. This is not about dumbing down. It is about the advanced practitioner taking responsibility for translation, recognising that the quality of the formulation is only as good as its ability to inform the responses of everyone who encounters the person.

Consider a concrete worked example. Amara is a twenty-six-year-old woman who has been self-harming by cutting for approximately eight years. She is currently known to a community mental health team, where she has a care coordinator and sees a psychiatrist quarterly. She also attends her GP surgery regularly for wound care, and she has recently begun working with a voluntary sector domestic abuse service following the end of a coercive relationship. Her sister is closely involved in her day-to-day life and is aware of the self-harm, though Amara has mixed feelings about how much her sister should know at any given time. The community mental health team holds a detailed formulation that understands Amara's self-harm as a long-standing affect regulation strategy rooted in early experiences of emotional invalidation, exacerbated by the chronic hyperarousal and dissociation associated with her trauma history, and currently being tested by the destabilising aftermath of leaving the abusive relationship. The GP, however, has no access to this formulation. The practice nurse who dresses Amara's wounds has developed her own informal understanding, which is broadly sympathetic but focused primarily on wound severity as a proxy for risk level. The domestic abuse worker is skilled in her own field but has limited training in self-harm and is uncertain how to respond when Amara discloses that she has been cutting more frequently since leaving the

relationship.

In this scenario, the advanced practitioner working within the community mental health team has a responsibility that extends beyond the boundaries of their own service. The formulation they hold needs to be shared, with Amara's consent, in a form that is useful to the GP, the practice nurse, and the domestic abuse worker. This does not mean sending a lengthy clinical letter that will be filed and never read. It means thinking carefully about what each professional needs to know in order to respond helpfully rather than harmfully. The practice nurse needs to understand that wound severity is not a reliable indicator of psychological distress in Amara's case, and that a calm presentation during wound care does not mean the risk has reduced. The domestic abuse worker needs to understand that an increase in cutting frequency following the end of the relationship is consistent with the formulation and does not necessarily indicate a crisis requiring immediate escalation, but that certain specific changes in the pattern of self-harm would warrant urgent contact with the care coordinator. The GP needs to know that Amara's self-harm is being actively addressed within a trauma-informed framework and that the most helpful response during consultations is one that acknowledges the self-harm without catastrophising or withdrawing, rather than defaulting to emergency referral each time wounds are presented.

This kind of coordinated communication requires the advanced practitioner to invest time in relationships with professionals outside their own service, and to resist the pull towards insularity that can develop in specialist settings. It also requires a clear agreement with Amara about what is shared, with whom, and in what form. The principle here is not simply procedural. It reflects a fundamental ethical commitment to the person's autonomy and to the therapeutic relationship as the foundation on which all safety planning rests. If Amara discovers that information has been shared without her knowledge or in ways she did not anticipate, the damage to trust may be irreparable, and the safety planning that depends on that trust will be undermined.

Family and support network involvement requires the same careful, person-centred approach. The default in many services has historically been to involve family members as a matter of course, on the assumption that more eyes watching and more people available to respond must be safer.

This assumption does not hold up under scrutiny. Family members may themselves be sources of stress, conflict, or even harm. The dynamics within a family system may be directly relevant to the formulation of why the self-harm is occurring. A family member who is told about a safety plan without understanding its rationale may respond in ways that are counterproductive, such as removing all sharp objects from the home in a manner that feels controlling and humiliating rather than caring, or monitoring the person so closely that the self-harm goes underground and becomes harder to address. Involvement of family or support network members should always be led by the person themselves, reflecting what they want and what they believe will be helpful, rather than being imposed as a professional default.

Where a person does want family involvement, the advanced practitioner's role is to help that involvement be genuinely useful. This means meeting with the family member or members, explaining the formulation in accessible terms, and helping them understand what helpful responses look like in practice. It means addressing the family member's own distress and fear, which are entirely understandable but which, if unaddressed, can lead to responses that inadvertently increase the person's distress. It means being honest about the limits of what any safety plan can guarantee, while also communicating clearly what the plan is designed to do and how the family member can contribute to it. In Amara's case, her sister's involvement needs to be negotiated carefully. Amara values her sister's support but is concerned that her sister becomes very anxious when she knows Amara has been cutting, and that this anxiety sometimes leads to arguments that leave Amara feeling worse. The advanced practitioner might work with Amara to think about what she would like her sister to know, what she would like her sister to do when she is struggling, and how to have that conversation with her sister in a way that feels manageable. The sister might also benefit from some direct psychoeducation about self-harm and about how to respond in a way that is supportive rather than escalating.

Across all of this work, the consistent thread is the formulation. A good formulation does not belong to one professional or one service. It belongs to the person, and it should be capable of informing every encounter that person has with every professional involved in their care. The advanced practitioner's task is to ensure that the formulation remains alive and coherent as it moves across settings, that it is updated when new

information emerges, and that contradictory narratives are identified and resolved rather than allowed to persist in parallel. This is demanding work. It requires clinical confidence, strong communication skills, a willingness to engage with professionals from very different backgrounds and with very different frameworks, and a sustained commitment to keeping the person's experience and perspective at the centre of every conversation. It is also, when done well, among the most protective things a practitioner can do for someone whose self-harm has become embedded in a complex, multi-system life.

In the case of Amara, the advanced practitioner would need to establish a regular review process that coordinates across all involved professionals. This might involve quarterly meetings with Amara, her sister, and representatives from each service, with the formulation serving as the anchor document. At each review, the team would assess whether the formulation remains accurate and whether the safety plan is still fit for purpose. For example, as Amara's relationship with her sister evolves, or as she develops new coping strategies, the formulation would need updating to reflect these changes. The advanced practitioner would then ensure these updates are communicated to all professionals in accessible formats, perhaps through a brief summary document that highlights key changes and their implications for each setting.

A common pitfall in this work is allowing formulations to become static documents rather than living frameworks that evolve with the person's circumstances. Practitioners may fall into the trap of treating the initial formulation as definitive, failing to recognise that people's lives, relationships, and coping mechanisms change over time. This is particularly problematic when working across multiple settings, as outdated formulations can lead to inconsistent responses and contradictory narratives. To avoid this, the advanced practitioner should build regular review points into the safety plan itself, asking Amara and others involved to reflect on what has changed and what remains helpful. Supervision provides an important space for examining these dynamics, with the practitioner bringing specific examples of how the formulation is being tested in real-world encounters and whether it continues to guide effective responses across different professional contexts.

This approach ensures that the formulation remains relevant across settings

and over time, preventing the emergence of contradictory narratives as circumstances change. It also maintains the person's agency at the centre of the process, with Amara actively participating in determining what information is shared and how her support network responds. The advanced practitioner's role is not to have all the answers, but to facilitate a process where the formulation continues to make sense of complex experiences and guide coordinated action across multiple systems, ultimately creating a safety net rather than a patchwork of disconnected interventions.

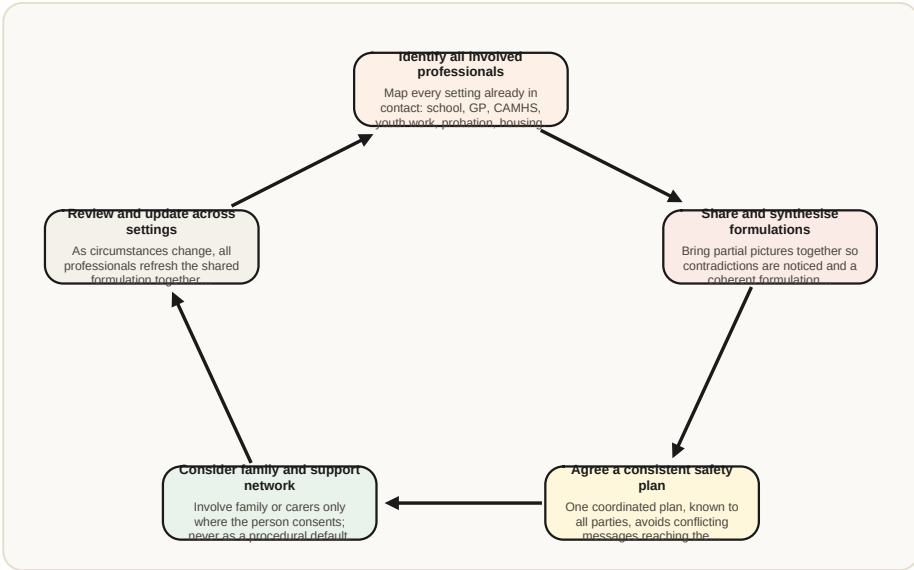


FIGURE 2.1 Coordinating Care Across Settings

Each stage builds a shared, consistent picture of the person who self-harms across all involved professionals and networks. Without completing the cycle, practitioners hold contradictory partial views that increase risk.

Points worth keeping

- Self-harm presents across settings beyond mental health services
- Family involvement where the person wants it, not as a default
- Coordinating formulation and safety planning across multiple professionals
- Keeping a consistent picture rather than contradictory ones

PRACTICE TASK

Outline how you would share the essentials of a person's safety plan with another professional involved in their care, while respecting what the person has consented to share.

CHAPTER 3

Transitions, Endings and Handovers in Complex Chronic Risk Work

When a clinician works with a person whose risk is chronic and recurrent, the therapeutic relationship itself becomes one of the most significant variables in the formulation. The relationship holds history, shared understanding, and — often — a carefully negotiated set of agreements about how crisis will be recognised and managed. It is precisely because that relationship carries so much clinical weight that the moment of transition, whether a planned ending, a transfer to another service, or an unplanned break, becomes one of the highest-risk periods in the entire course of care. Advanced practice in this area requires the clinician to think about transitions not as administrative events but as formulation problems in their own right, demanding the same rigour, collaborative spirit, and attention to individual meaning that characterises good risk formulation at any other point in the work.

The literature on therapeutic endings with high-risk populations is less developed than the literature on risk assessment itself, yet the clinical stakes are well established. Ougrin, Tranah, Stahl, Moran and Asarnow (2015), reviewing self-harm interventions in adolescents, noted that gains made during active treatment were vulnerable to erosion at the point of discharge or transfer, particularly where the therapeutic alliance had been the primary mechanism of change. Knox and Cohn (2010) drew attention to the way in which endings can reactivate attachment-related distress, and for individuals whose chronic risk is rooted in early relational trauma or insecure attachment — as is frequently the case in presentations involving emotionally unstable personality disorder, complex post-traumatic stress disorder, or long-standing depression with suicidality — the ending of therapy can function as a recapitulation of earlier experiences of abandonment or rejection. The clinician who fails to hold this in mind risks inadvertently creating the very conditions that the formulation identifies as precipitants for crisis.

Formulating the transition itself is therefore not optional. It means returning

explicitly to the longitudinal formulation that has been built over the course of the work and asking a specific set of questions: what does this ending mean to this person, given their history and the meanings they have constructed around loss, rejection, and care? What protective factors have been contingent on the therapeutic relationship, and which of those can be transferred to other relationships or internalised? What early warning signs are most likely to emerge in the period immediately following transition, and who will be positioned to notice them? These questions are not merely clinical housekeeping; they represent a distinct formulation task, and the answers should be documented with the same care as the original risk formulation.

The concept of a transferable formulation is central to good handover practice. A transferable formulation is not simply a summary of diagnosis and risk history; it is a narrative account of the person's risk in context, written in a way that enables a receiving clinician, team, or service to understand not only what has happened but why, and what conditions are likely to be associated with deterioration or improvement. Adshead and Sarkar (2012), writing about personality disorder and risk, argued that the most dangerous handovers are those in which risk is communicated as a static property of the person — a score, a category, a label — rather than as a dynamic process embedded in relational and contextual factors. A receiving clinician who is told only that a person is "high risk" or "has a history of serious self-harm" is poorly equipped to respond to that person's actual presentation; a clinician who understands that risk escalates when the person experiences perceived criticism from authority figures, that it is mitigated by structured daily routine and contact with a specific friend, and that the person's characteristic early warning sign is social withdrawal rather than overt distress, is in a position to act with genuine clinical intelligence.

Writing a transferable formulation requires the clinician to make explicit what has often remained implicit in the therapeutic relationship. Much of what an experienced clinician knows about a person with chronic risk has been built up through repeated contact, through noticing patterns across episodes, through the texture of the therapeutic relationship itself. The discipline of writing it down in a form that is useful to someone who has never met the person is considerable, and it is a discipline that should be practised in supervision rather than left to the final weeks of a piece of work. The formulation should include, at minimum: a brief account of the

developmental and relational context of the person's risk; a description of the cycle of risk as it has been understood collaboratively with the person; the specific precipitants, maintaining factors, and protective factors that have been identified; the person's own understanding of their risk and what helps; and a clear account of what the person finds most difficult about transitions and endings, including any previous experiences of service endings that have been associated with deterioration.

To illustrate what this looks like in practice, consider the following worked example. A woman in her mid-forties — referred to here as Diane — has been seen in weekly individual therapy for two years following a serious suicide attempt. Her formulation identifies a pattern rooted in early experiences of emotional neglect and an abusive first marriage: she experiences intense shame and a sense of being fundamentally unlovable when she perceives herself to have failed or disappointed others, and this shame state is the proximate driver of suicidal ideation. Over the course of therapy she has developed a more nuanced understanding of this pattern, has built a small but meaningful social network, and has engaged consistently with a safety plan that includes contact with her sister as a first-line protective step. The therapy is ending because the service has a time-limited contract and the funding period has concluded. Diane has expressed ambivalence about the ending, oscillating between apparent acceptance and episodes of increased hopelessness in sessions. The clinician formulates the ending itself as a potential precipitant: the ending of therapy carries the meaning, for Diane, of having been found wanting — of having failed to get better quickly enough, of being too much trouble, of being abandoned by someone she trusted. The clinician works explicitly with this meaning in the final sessions, not to resolve it entirely but to name it, to locate it within Diane's broader formulation, and to help Diane develop a narrative in which the ending is understood as a structural feature of the service rather than a verdict on her worth. The transferable formulation written for the receiving community mental health team includes a specific section on the meaning of endings for Diane, notes that the period immediately following transition is likely to be the highest-risk period, recommends that the first contact with the new worker be offered within two weeks rather than the standard four, and identifies the sister as the key person to be kept informed. The safety plan is reviewed and updated to include a specific step for the transition period — a scheduled phone call with the sister on the day of the final session — and a

copy is given both to Diane and to the receiving team.

This example illustrates several principles that apply more broadly. First, the formulation of the transition is continuous with, not separate from, the formulation of the risk itself. The meaning of the ending is not a separate psychological event; it is the formulation's predictions playing out in real time. Second, the safety plan must be updated to reflect the specific conditions of the transition period. A safety plan written for ongoing therapy may not be adequate for the period of handover, when the primary protective relationship is temporarily absent and the person's distress is likely to be elevated. Third, communication with the receiving service must be specific and formulation-based. General risk warnings without contextual explanation are not only clinically inadequate; they can actively harm the therapeutic relationship with the new clinician by priming them to relate to the person primarily through the lens of risk rather than through the lens of understanding.

The question of timing is also a formulation question. Research on the sequencing of therapeutic endings with high-risk populations consistently suggests that abrupt endings are associated with worse outcomes than planned, gradual endings with adequate preparation time. Guthrie and colleagues (2001), in their work on brief psychodynamic interpersonal therapy for deliberate self-harm, noted that the therapeutic frame around ending — including explicit discussion of the ending's meaning and explicit preparation for the post-therapy period — was a significant component of the intervention's effectiveness. In complex chronic risk work, this means that the ending should be introduced into the formulation and the safety planning process well in advance of the final session, not raised for the first time in the penultimate meeting. Where service structures make this difficult — where funding decisions are made at short notice, or where a clinician's departure is unplanned — the clinician has a professional responsibility to advocate for adequate transition time and to document clearly when that advocacy has been unsuccessful.

Unplanned endings present particular challenges. When a clinician leaves a service unexpectedly, when a person disengages from therapy without notice, or when a crisis leads to an unplanned admission that disrupts the therapeutic relationship, the formulation must be revisited urgently. The person's response to the unplanned ending is itself clinically significant data:

it will often confirm or refine the formulation's account of the person's relational vulnerabilities. A person who responds to an unplanned ending with immediate escalation of risk is communicating something important about the degree to which the therapeutic relationship has been functioning as a protective factor, and this information should be incorporated into the formulation that is passed to whoever takes over the care. The receiving clinician should be supported, through supervision and through the quality of the documentation they receive, to understand that the elevated risk at the point of handover is a predictable consequence of the transition, not a sign that the previous work was ineffective or that the person is simply deteriorating.

There is also an important dimension of power and voice in transition work that advanced practitioners must hold in mind. People with chronic risk histories frequently have extensive experience of services, and that experience is often marked by repeated endings, transfers, and handovers, many of which have been experienced as abrupt, unexplained, or disrespectful. The cumulative effect of multiple poor transitions can itself become a formulation factor: a person who has been handed over between services many times without adequate explanation or preparation may have developed a well-founded expectation that care relationships are unreliable, and this expectation will shape how they engage with any new clinician. Acknowledging this history explicitly, validating the person's experience of previous transitions, and demonstrating through the quality of the current transition that a different experience is possible, is not merely good manners; it is a clinical intervention in its own right, one that has the potential to modify a formulation factor that has contributed to chronic risk over many years.

Finally, the clinician's own emotional response to ending with a high-risk person deserves attention in supervision. The anxiety of handing over someone whose safety has felt, in some measure, like a personal responsibility is real and clinically significant. Clinicians who are not supported to process this anxiety may either prolong endings beyond what is clinically indicated — inadvertently reinforcing the person's dependence on the therapeutic relationship — or may manage their own distress by becoming prematurely detached, communicating to the person that the relationship was less significant than it appeared. Neither response serves the person well. Supervision in the period around a complex transition

should attend explicitly to the clinician's experience of the ending, and the supervisor should be alert to signs that the clinician's own attachment processes are influencing the clinical decisions being made about timing, communication, and the content of the transferable formulation.

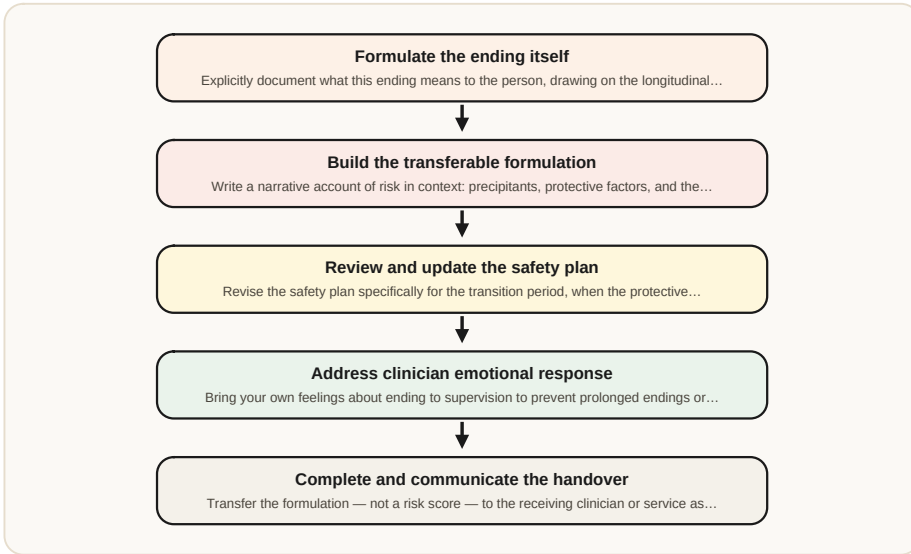


FIGURE 3.1 Managing Transitions in Complex Chronic Risk Work
Each stage must be completed before the next; the sequence shows that a transition is a clinical process, not an administrative event. The transferable formulation is the central product that carries meaning, risk context, and protective factors across the handover.

Points worth keeping

- Transitions and endings in complex chronic risk work are formulation problems in their own right, not administrative events: the meaning of the ending to the person must be explicitly formulated and documented, drawing on the longitudinal case formulation.
- A transferable formulation — a narrative account of the person's risk in context, including precipitants, protective factors, and the meaning of endings — is the essential clinical product of a well-managed handover, and is more clinically useful than a risk category or score.
- Safety plans must be reviewed and updated specifically for the transition period, when the primary protective relationship is temporarily absent and risk of deterioration is typically at its highest.

- The clinician's own emotional response to ending with a high-risk person is clinically significant and should be addressed explicitly in supervision, to prevent either prolonged endings or premature detachment from distorting the transition process.

PRACTICE TASK

Select a current or recent case involving a person with chronic or recurrent risk where a transition, ending, or handover is approaching or has recently occurred. Draft a transferable formulation for that person, structured as a narrative rather than a checklist, which includes: a brief account of the developmental and relational context of their risk; a description of the cycle of risk as you have understood it collaboratively with the person; the specific precipitants, maintaining factors, and protective factors you have identified; the person's own understanding of their risk and what helps; and a specific section on what this person finds most difficult about transitions and endings, including any previous experiences of service endings associated with deterioration. Then review the person's current safety plan and identify any steps that need to be added or modified specifically for the transition period. Bring both documents to supervision and use them as the basis for a discussion about the meaning of the ending in the context of the formulation, and about your own emotional response to the transition.

CHAPTER 4

Distinguishing chronicity from acuity: formulating risk when crisis becomes a recurring pattern rather than a single episode

Most training in risk formulation and safety planning assumes a single index episode: a person presents in crisis, a formulation is built around that episode's circumstances, and a safety plan follows from it. That model holds up reasonably well for a first presentation. It becomes far less useful once a person has been in and out of crisis for months or years, because the central question changes. The question is no longer "what has brought this person to this point", but "what does it mean that this person has been at a broadly similar point many times before, and what, if anything, is actually different this time." Practitioners trained only on the single-episode model tend to default to one of two responses when faced with a chronic or recurrent pattern: they either reuse the original formulation indefinitely, treating every subsequent crisis as a repetition of the first, or they build a fresh formulation from scratch each time, as though the person's history carried no information at all. Neither is adequate. This lesson sets out a third approach: treating the chronic pattern itself as the object of formulation, with its own baseline, its own warning signs, and its own rules for what counts as a genuine deviation worth acting on differently. Getting this distinction wrong has real consequences in both directions: a formulation that never updates risks missing exactly the deterioration it exists to catch, while a formulation rebuilt from nothing each time wastes the one genuine advantage a chronic presentation offers over a first one, namely a documented history to reason from.

The practical difference between formulating an acute presentation and a chronic one is mostly a difference in what the practitioner is comparing the person's current state against. In an acute, first-presentation formulation, the comparison is implicit: current state against an assumed ordinary baseline the practitioner has to infer, because there is no documented history to check against. In a chronic presentation, the comparison should be explicit:

current state against this specific person's own, previously documented pattern. That shift changes what a "trigger" means. For a first presentation, almost anything notable in the recent history might be the trigger, because the practitioner has no basis for ruling anything out. For a chronic presentation, the useful triggers are the ones that explain why today is different from the person's usual run of difficult days, not the ones that would explain why they are having a difficult day at all. A chronically high level of background distress is not, on its own, informative; a change in its shape, frequency, or the person's own relationship to it, is. Learning to ask "what is different about now, compared with their own baseline" rather than simply "what is wrong" is the central shift this lesson asks practitioners to make.

NICE NG225 is explicit that risk-scoring tools and checklists should not be used to predict future self-harm or suicide, because the factors they weight are derived from population-level statistical association rather than individual-level causal mechanism. This matters more, not less, in chronic presentations, for a specific reason. With a single-episode presentation, practitioners are at least cautious about over-trusting a checklist, because they know they are working with limited information. With a chronic presentation, the temptation runs the other way: a long history of crises that have not resulted in death or serious injury starts to look, informally, like its own predictive evidence, feeding a quiet assumption that "they always do this and they're always alright." That assumption is exactly the population-level error NG225 warns against, applied to a sample size of one person's own history rather than to a population, but no more valid for it. A given person's survival of twenty previous crises says nothing certain about the twenty-first, and formulation work with chronic presentations has to hold that explicitly, in documented reasoning, rather than let it operate as an unexamined background assumption. The corrective is not to treat every recurrence as acutely as the first one; it is to make the reasoning behind treating it differently, or the same, visible and checkable, which is what the baseline formulation below is for.

A workable baseline formulation for a chronic or recurrent presentation should capture, in writing, four things: the person's usual pattern of urges, self-harm, or crisis, including its typical frequency, intensity, duration, and how it usually resolves; their usual early warning signs, described specifically enough to be checked against rather than left as generic descriptors; their

usual coping repertoire, including what has helped before and what has not; and an explicit statement of what would count as a genuine change from that pattern, agreed where possible with the person themselves. The fourth element is the one most often missing from existing notes, and it is the one that does the actual work. Without it, every subsequent review has to re-derive, informally and inconsistently, what would count as worse, which is exactly the condition under which both of the failure modes described below take root, because different practitioners silently apply different unwritten thresholds. With it, a new practitioner picking up the case for the first time, an out-of-hours service, or a supervisor doing a spot check, all have the same explicit reference point for the question “is this the same pattern or has it changed”, rather than relying on one clinician’s accumulated, undocumented sense of the case.

A composite, non-identifying example illustrates how this works in practice. Consider a person with a four-year history of self-harm connected to conflict with a partner, presenting roughly monthly, with corresponding contact with services each time, whose baseline formulation documents this pattern explicitly: frequency, usual triggers, usual method of self-soothing afterwards, and a working statement that self-harm connected to relationship conflict, resolving within the same week, represents the known pattern rather than escalation. A review six months later records a presentation that does not fit this description: the self-harm follows a job loss, not relationship conflict, is more severe than the person’s own documented typical range, and is accompanied, for the first time, by the person describing having thought the situation through in advance rather than acted on impulsively. None of these three features would be especially dramatic, taken alone, in an acute, first-presentation assessment, but set against the documented baseline, together they constitute exactly the kind of deviation the baseline formulation was built to surface: a different trigger category, a severity outside the known range, and a marker of forward planning that had not previously been part of the pattern. The review correctly treats this as a genuine change requiring a fresh safety plan rather than a repetition of the known pattern, and the written baseline is what allows that judgement to be made on specific, checkable grounds rather than on a feeling that “something is different this time” that a less experienced colleague, or the same colleague on a difficult day, might not have picked up on.

Baseline is not a synonym for mild, and none of the above is an argument for treating a person's habitual severity as acceptable simply because it is familiar. Some people's genuine, long-standing pattern already includes serious harm, previous hospital admissions, or previous near-fatal episodes, and the task set out here is still to compare the current presentation against that specific person's own pattern, not against an external severity scale that would label every instance of it an emergency. Confusing baseline with mild produces exactly the alarm fatigue described below; confusing baseline with acceptable produces a different, quieter failure, in which a service stops actively safety-planning around a known severe pattern simply because nothing about it currently looks new. A baseline that includes serious harm still warrants an active, current safety plan, supervision oversight, and the same explicit review cycle as any other case; what changes is only the question asked at each presentation, which is whether this instance sits inside that known, serious pattern or outside it, not whether the pattern itself is serious enough to warrant concern.

Chronic risk work fails in two directions, and NG225-aligned practice has to guard against both rather than treating caution against one as sufficient protection against the other. The first failure mode is alarm fatigue or complacency: familiarity with a recurring pattern breeds an assumption that this presentation, like the last several, will resolve the same way, leading practitioners to under-respond to genuine deviations because they have stopped actively checking the current presentation against the documented baseline at all. The second, less often discussed, failure mode is defensive overreaction: a practitioner who has been anxious about missing a genuine deterioration starts treating every instance of a person's ordinary, already-documented baseline distress as a fresh crisis requiring escalation, inpatient referral, or intensified contact that the person's own history does not support, and that can itself be experienced as punitive or infantilising, eroding the therapeutic relationship the longer-term work depends on. Both failure modes share the same root cause: the absence of an explicit, written standard for what counts as a change, which forces every individual practitioner's judgement, on every occasion, to substitute for a documented decision the service has already made once, carefully, and can therefore apply consistently.

Because the baseline formulation is doing real clinical work, rather than functioning as a one-off assessment filed and forgotten, it needs version

control in the ordinary sense. Each formulation should be dated, should state explicitly whether it supersedes a previous version or merely adds detail to it, and superseded versions should remain accessible rather than being overwritten, so a reviewer can see not just the current baseline but how and when the understanding of the person's pattern has changed over time. A formulation that has not been reviewed within a defined period should trigger a scheduled review rather than persisting by default; services should set that period according to the frequency of the person's own pattern rather than applying one fixed interval across a whole caseload. In practice this means the formulation document itself should carry a last-reviewed date and a next-review trigger, in the same way a care plan would, rather than living as an undated narrative note that nobody can tell, six months on, whether anyone has actually reconsidered.

Supervision is the structural safeguard against both failure modes, but only if it is used to check the specific question this lesson has set out, rather than functioning as a general update. A supervision discussion of a chronic case should routinely ask three things: what does the documented baseline currently say, has anything in the recent presentation actually departed from it, and is the practitioner's current response proportionate to that departure, or simply to the fact that the person has presented again. Asking these questions does two things at once. It surfaces complacency, because a practitioner who has stopped checking against the baseline usually cannot answer the first question cleanly without checking the notes in the room, which is itself diagnostic. And it surfaces overreaction, because a practitioner escalating a presentation the baseline formulation already describes as the known pattern will struggle to articulate what specifically has changed. A later module in this course returns to supervision's role at the organisational level; the point to take from this lesson is narrower and more immediate: chronicity is not a reason to formulate less carefully than a first presentation, but a reason to formulate against an explicit, written, regularly reviewed standard that makes departures from the person's own pattern visible to whoever is looking, not just to whoever has carried the case the longest.

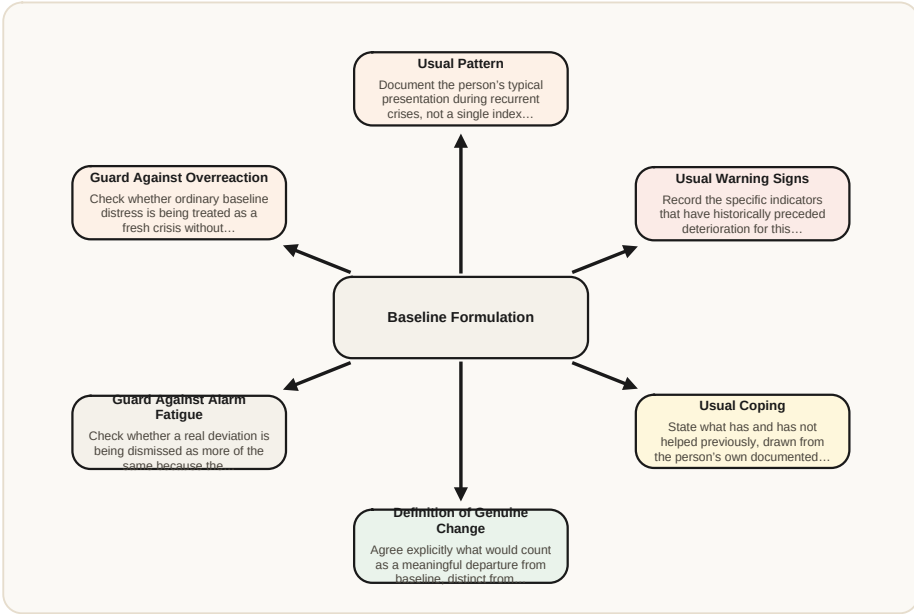


FIGURE 4.1 Formulating Risk in Chronic and Recurrent Presentations

The baseline formulation is the central reference point against which all subsequent crisis episodes are compared. Both failure modes — alarm fatigue and defensive overreaction — arise when practitioners lose sight of the documented baseline.

Points worth keeping

- Formulate a chronic or recurrent presentation against the person's own documented baseline, not against a fresh, implicit assumption each time
- A baseline formulation should state usual pattern, usual warning signs, usual coping, and an explicit, agreed definition of what would count as a genuine change
- Guard against both failure modes: alarm fatigue (treating a genuine deviation as more of the same) and defensive overreaction (treating ordinary baseline distress as a fresh crisis)
- Date and version every formulation, and use supervision to check specifically whether the current response is proportionate to a documented departure from baseline

PRACTICE TASK

Take a case from your own caseload with a recurrent pattern of risk and draft its baseline formulation using the four elements set out in this lesson. Then identify one thing that would have to be true of a future presentation for you to treat it as a genuine change rather than the known pattern.

CHAPTER 5

A functional analysis of recurrent self-harm: understanding what the pattern does for the person, not just what it risks

A risk assessment asks how likely future self-harm is and how severe it might be. A functional analysis asks a different question: what does this particular act of self-harm actually do for this particular person, in this particular situation. The two questions are not interchangeable, and treating them as though they were produces a specific, recurring error in chronic risk work. A service that attends only to frequency and severity, without asking what the behaviour is for, tends to respond to every episode with the same escalating menu of containment measures: closer monitoring, more frequent contact, removal of means, admission. Where self-harm serves a function the person has not found another way to meet, none of these measures addresses the underlying problem; they can suppress the specific behaviour being watched for while leaving the person to meet the same need some other way, occasionally one that is harder to detect and respond to than the original. Functional analysis is the tool that moves formulation from describing a pattern, which was the subject of the previous lesson, to explaining it, which is what any intervention beyond containment has to be built on. This lesson sets out how to conduct a functional analysis with someone who self-harms repeatedly, the functional categories most often encountered in this work, and how the resulting formulation should shape safety planning rather than sit alongside it as a separate piece of paperwork.

The basic structure borrowed from behavioural analysis is antecedent, behaviour, consequence, and it earns its place here because it answers the question that matters clinically: why does a behaviour with serious longer-term costs persist. The antecedent is whatever internal state or external event precedes the act, which in practice is usually an internal state, such as a specific, often hard-to-name quality of distress, numbness, or agitation, sometimes triggered by an identifiable external event and sometimes arising without one the person can point to. The behaviour itself

is the self-harm, and a formulation needs enough detail about its usual form to notice when that form changes, without the notes becoming a procedural account that serves no clinical purpose. The consequence is where the explanatory power sits, and it has to be examined on two timescales at once: the immediate consequence, within seconds or minutes, and the longer-term consequence, over hours, days, or the life of the relationship with services. Self-harm persists, in the great majority of cases, because its immediate consequence is some form of relief, a drop in unbearable arousal, a return of a sense of control, or a feeling of having finally been heard, even while its longer-term consequences are physical harm, shame, and strain on relationships and on contact with services. A formulation that only records the longer-term costs, which are usually obvious to everyone involved including the person themselves, without identifying the immediate consequence that is actually doing the maintaining, has not explained anything; it has simply restated that the behaviour is harmful, which nobody disputed.

Several broad functional categories recur across the research and clinical literature on self-harm, and it helps to hold them loosely rather than as a fixed checklist, because a single episode can serve more than one function at once, and the same person's function can shift between episodes or over the course of treatment. Affect regulation, bringing down an intolerable level of internal arousal or distress, is the most consistently reported function and the one most directly tied to the immediate-relief mechanism described above. A communicative function, conveying distress that words have not successfully conveyed, or that the person does not feel entitled or able to voice directly, is common and is often the function most readily dismissed by services, for reasons addressed below. Self-punishment, directed at the self in response to shame, self-criticism, or a sense of having done wrong, is another recognised function, distinct from affect regulation in its underlying logic even though it can look similar from outside. A fourth, less discussed but clinically important function is restoring a sense of agency or control in circumstances that otherwise feel completely unmanageable. None of these functions is more legitimate or more concerning than another in itself; what matters clinically is identifying which one, or which combination, is actually operating for this person, in this presentation, rather than assuming a single universal explanation applies.

A brief composite example shows how more than one function can be at

work in what looks, from outside, like a single repeated behaviour. Consider a person whose self-harm has, for the first two years of contact with services, consistently followed episodes of acute anxiety and dissociation, with an immediate consequence, each time, of feeling returned to her body and able to think clearly again; the functional analysis across this period points clearly to affect regulation. Eighteen months into treatment, a new pattern emerges following disputes with a partner who has previously minimised her distress, where the immediate consequence the person describes is not relief from dissociation but her partner's visible alarm and a change in his behaviour towards her for several days afterwards. The behaviour looks identical on the surface; the function has shifted to something closer to communication, and a safety plan still built around dissociation-focused grounding, which worked well for the earlier pattern, will do little for the new one, because it addresses a function that is no longer the one operating. Recognising the shift required asking, specifically, about each recent episode rather than relying on the formulation already on file, which is the practical discipline this lesson is really asking for.

This points to a wider principle: a functional formulation is a hypothesis about a specific period of a person's life, not a permanent fact about them, and it should be revisited with the same explicit rigour as the baseline formulation in the previous lesson, rather than carried forward unexamined once it has been written down. A functional analysis completed a year ago and never revisited risks becoming exactly the kind of comfortable, unexamined assumption this course has already warned against. Building a brief re-check into existing formulation reviews, simply asking whether the immediate consequence described still matches what the person reports now, keeps the functional account doing genuine explanatory work rather than ageing quietly into received wisdom about the case, and the updated account should be shared with whoever else is coordinating the person's care, for the same reasons set out in the lesson on working across settings.

"Attention-seeking" deserves specific attention because it remains common informal shorthand in services, and it causes real harm when it substitutes for an actual functional analysis. Stripped of its judgemental framing, what the label is usually gesturing at is a communicative function: the person is trying, through the self-harm, to secure a response, a change, or acknowledgement that other channels have not produced. Dismissing that as manipulative rather than recognising it as communicative tends to

escalate rather than resolve the underlying need, for the straightforward reason that a person whose attempt to be heard is met with punishment or withdrawal usually either disengages from services or escalates the behaviour in an attempt to be heard more forcefully, neither of which is a safer outcome than taking the communicative function seriously in the first place. Taking it seriously does not mean treating every request as one that must be granted; it means formulating what is actually being communicated, and building a response to that, rather than to the behaviour in isolation. A communicative function is a legitimate function to understand and respond to, in the same way any other function is, and naming it accurately, as communication rather than manipulation, changes what a sensible response looks like.

Conducting the functional analysis itself works best as a collaborative, curious conversation anchored to one recent, specific incident rather than a general discussion of “why you self-harm”, which tends to produce rehearsed or overly general answers. Useful questions explore what was happening and what the person was feeling in the period immediately before the incident, what, specifically, changed or felt different immediately afterwards, and what the person believes would have needed to happen instead for the urge to pass without the behaviour. Timing matters: a conversation attempted in the acute aftermath of an incident often meets a person who is not yet able to reflect on it usefully, while one left too long risks losing accurate recall of what actually preceded and followed it, so a window of a few days, once the immediate crisis has settled but while the incident is still freshly recalled, usually works best. The tone matters as much as the questions: a functional analysis conducted as an interrogation looking for the “real reason” produces guarded, defensive answers, while one conducted with genuine curiosity about what the behaviour achieves, stated plainly as an attempt to understand rather than to judge, tends to produce the detail the formulation actually needs.

The functional formulation should then shape safety planning directly, and this is where the work pays off. Once a function is identified, such as bringing down unbearable arousal, the task is not to hand the person a generic list of distraction techniques, which research and clinical experience both suggest tend to fail precisely because they ignore function and are applied uniformly regardless of what any individual episode is actually for. The task is to work collaboratively with the person to identify what else,

specific to them and to that function, could meet the same underlying need, built from what has genuinely helped them before, rather than from a standard list handed out regardless of formulation. For a communicative function, the equivalent work is agreeing, in advance and with the person, a different, more direct way for them to signal that level of distress to someone who can respond, so the self-harm is no longer the only available channel. This is slower and more individualised than issuing a standard coping-skills handout, and it is also considerably more likely to actually reduce the behaviour, because it addresses what is maintaining it rather than only what is observable about it.

The resulting formulation belongs in the case notes and in the safety plan's own stated rationale, not filed separately as background psychological material, because a plan built without reference to function is a plan built on guesswork about why it should work. It should also be shared, in formulation terms rather than as a diagnostic label, with anyone else coordinating the person's care, so that a different practitioner responding to a future presentation is working from the same explanatory account rather than reconstructing one from scratch under pressure. Handled this way, functional analysis is what gives the chronic-pattern baseline described in the previous lesson its explanatory depth: knowing that a pattern repeats is a necessary starting point, but knowing what the pattern does for the person is what any intervention beyond simple containment actually has to be built on.

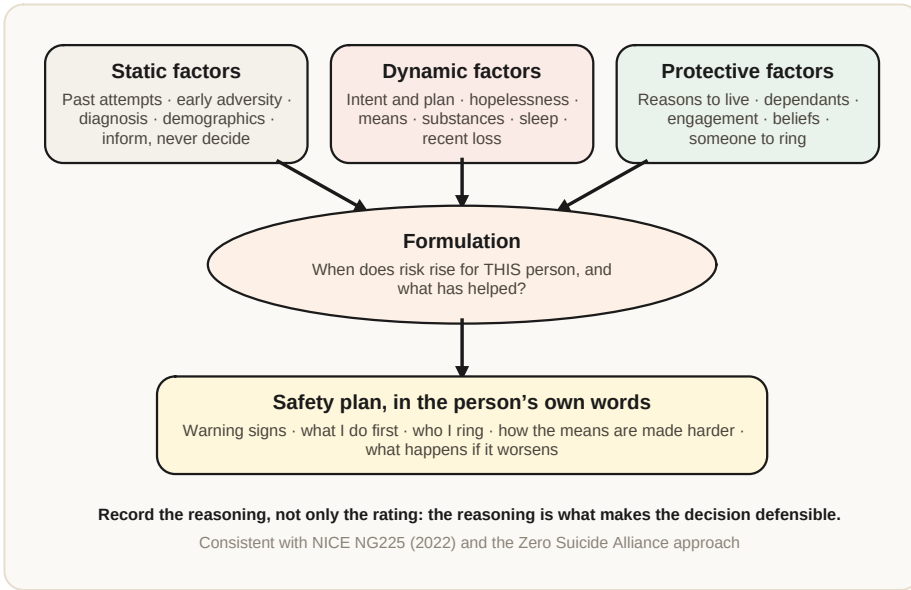


FIGURE 5.1 From risk assessment to risk formulation

A defensible approach to risk does not predict, it formulates: it separates the factors that cannot change from those that can, names what protects, and converts the picture into a plan somebody could actually follow at three in the morning. A checklist score recorded without a formulation and a plan is a record of an assessment, not the management of a risk.

Points worth keeping

- A functional analysis asks what the self-harm does for the person, not just how risky or frequent it is, using an antecedent-behaviour-consequence structure
- Common functions include affect regulation, communication, self-punishment, and restoring a sense of control; a single episode can serve more than one, and this can shift over time
- Treat a communicative function as communication to be understood and responded to, not dismissed as manipulation, which tends to escalate rather than resolve it
- Build safety planning around the identified function collaboratively with the person, rather than issuing a generic, uniform coping-skills list that ignores what the behaviour is actually for

PRACTICE TASK

Using a recent incident from a case you know well, draft the antecedent, behaviour and immediate-consequence sequence as the person themselves would likely describe it, then identify which functional category or categories this points towards.

PART 2

Sustaining Collaborative Practice in Long-Term Risk Work

This module addresses the clinical and relational challenges of maintaining effective, recovery-oriented safety planning over time, including complacency, diagnostic labelling and revising plans that repeatedly break down.

CHAPTER 6

Avoiding alarm fatigue and complacency over a long course of chronic risk work

Alarm fatigue in chronic risk work is a predictable, well-recognised phenomenon, not a personal or moral failing in the individual practitioners it affects, and that distinction matters practically as well as in principle. A culture that treats drift as evidence of carelessness or poor character makes the problem harder to catch, not easier, because practitioners who notice it happening in themselves have every incentive to hide it rather than disclose it, for fear of exactly the judgement that framing invites. A culture that instead treats alarm fatigue as a known occupational hazard of this specific kind of work, comparable to the habituation effects documented in other high-repetition monitoring roles, makes it something a practitioner can raise in supervision as a current risk to manage, rather than a failure to confess. This lesson treats it as the former throughout, and sets out what drives it, how to recognise it early, what services can build structurally to guard against it, and why the opposite failure, driven by a related but distinct

mechanism, deserves equal attention rather than being treated as the safer error.

The underlying mechanism is a form of habituation that is adaptive in many settings and poorly suited to this one. Repeated exposure to the same category of crisis, without correspondingly severe outcomes, recalibrates a practitioner's internalised sense of what counts as alarming, in much the same way a person who works around loud machinery stops consciously registering the noise. In most settings this recalibration is simply efficient: it frees attention for genuinely novel information. In chronic risk work it is actively dangerous, because the stakes attached to a missed deterioration remain exactly as high on the fortieth presentation as they were on the first, even though the practitioner's felt sense of alarm has, entirely normally, diminished. The mismatch between unchanged stakes and diminished felt alarm is the whole problem, and no amount of experience or competence makes a practitioner immune to it; if anything, longer experience with a specific case increases exposure to the habituating pattern.

Drift is detectable earlier than it is usually caught, if a service knows what to look for rather than waiting for a practitioner to self-report concern. Clinical notes are often the first place it shows: entries that once described a presentation in specific, checkable detail compress, over months, into shorthand that effectively just restates "same as usual" without the specifics that would let a reader judge whether that is actually true. Assessment conversations shorten in a similar way, with fewer follow-up questions asked about a presentation that would have attracted closer enquiry earlier in the case. Supervision discussions of the case become shorter and more procedural, often moving on quickly because both supervisor and supervisee share an unspoken sense that this case is settled. And practitioners can often name the felt experience themselves, if asked directly rather than generally: a sense of "here we go again" arriving before the details of a new presentation have actually been reviewed. None of these signs, alone, proves drift has set in, but a pattern across several of them, sustained over a run of contacts, is a reliable enough signal to warrant the structured response set out below.

Several organisational factors make drift substantially more likely, and recognising them matters because they are the targets a service can actually act on, as distinct from an individual practitioner's attentiveness, which is

harder to intervene on directly. A caseload with limited variety, where several cases share a similar chronic pattern, accelerates habituation faster than a mixed caseload does, because the recalibration described above compounds across similar cases rather than resetting between them. Low supervision frequency relative to a case's complexity leaves drift unchecked for longer between the points at which it would otherwise be caught. Continuity without any planned opportunity for a fresh perspective, where the same practitioner holds a complex chronic case for years without a scheduled second look from anyone else, removes the single most reliable external check on drift. And a service culture that implicitly rewards calm, steady responses to repeat presentations while treating raised concern as overreaction, even unintentionally, teaches practitioners that under-responding is the socially safer error, which is precisely backwards given the mismatch described above.

Alongside the structural safeguards a service builds, there is a simpler practice an individual practitioner can use between them: a brief, periodic self-check, done honestly and privately rather than as a tick-box exercise, asking three plain questions about a specific long-standing case. Can I describe, specifically, what happened at the last two contacts, or would I struggle to do so without checking the notes. Have I asked a genuinely new question about this case in the last month, or have my questions become routine. And would I notice if this presentation changed, or am I confident mainly because nothing dramatic has happened recently. A practitioner who answers these honestly and finds the answers uncomfortable has caught something worth raising in supervision before a structural safeguard catches it instead; the self-check is not a substitute for the safeguards below, since by definition it depends on the very noticing that habituation erodes, but it costs nothing to run and occasionally works.

The structural safeguards that work are the ones that do not depend on a practitioner noticing their own drift, since the nature of drift is that it is hard to notice from inside it. A scheduled formulation refresh, at an interval set by the case's own pattern rather than a single fixed point across the whole caseload, forces the specific comparison against baseline set out in an earlier lesson, rather than leaving that comparison to happen only when someone remembers to make it. A mandatory second-opinion or peer-review point, built into the care pathway at defined intervals regardless of whether the regular practitioner has raised any concern, brings a genuinely fresh

perspective to a case precisely because the reviewer has not been exposed to the habituating repetition the regular practitioner has. Supervision that asks the three specific questions set out previously, rather than functioning as a general update, closes the loop between what the case notes say and what the practitioner's current response actually is. None of these measures depends on anyone admitting, in the moment, that they might have drifted, which is what makes them more reliable than good intentions alone.

The opposite failure deserves equal attention, and it is driven by a related but distinct mechanism worth naming separately. Where alarm fatigue comes from habituation, the opposite failure, defensive overreaction, often comes from a different source: vicarious stress or a specific anxiety about having missed something, sometimes following a difficult incident elsewhere in a practitioner's career or a service's recent history, that is not actually about the case in front of them. A practitioner driven by this mechanism needs a different response from a supervisor than a practitioner drifting through habituation does. Where habituation calls for structured re-engagement with the specific detail of the current case, anxiety-driven overreaction usually calls for support and space to process what is actually driving the heightened response, because escalating the care plan further will not resolve an anxiety that was never really about this particular presentation. Distinguishing which mechanism is operating, rather than assuming every departure from a measured response is the same problem, is itself a supervision skill.

A brief composite example shows a structural safeguard working as intended. A practitioner holding a complex chronic case for three years had, by the final year, reduced each monthly review to a short note confirming no change; a scheduled second-opinion review, built into the pathway regardless of concern, was conducted by a colleague with no prior exposure to the case, who asked the person directly about the past month and recorded a shift in the pattern the regular practitioner's notes had not captured, because the specific questions that would have surfaced it had quietly stopped being asked. Nothing about the regular practitioner's competence had changed; what had changed was exposure, and the safeguard caught what individual vigilance alone had, entirely predictably, stopped catching.

Team culture matters here in a way that is easy to overlook. A team in which

naming fatigue with a specific case, in an ordinary team meeting rather than only in formal supervision, is treated as a normal and useful contribution rather than as an admission of failing, catches drift earlier than a team where the only sanctioned place to raise it is a structured review months apart. Peer discussion, ordinary and informal, often surfaces exactly the kind of fresh-eyes observation a scheduled second-opinion review is designed to force, at a fraction of the cost, simply because colleagues who are not habituated to a particular case can hear a brief description of it and notice what the practitioner carrying it no longer does. Building this into a team's ordinary rhythm, rather than leaving it to chance or to practitioners who happen to be unusually open about struggling, is itself a protective structure worth deliberately cultivating.

None of this needs heavy documentation to be effective, but it benefits from a light trace. A line in a supervision record noting that a drift self-check was discussed, or that a peer raised a fresh observation about a long-standing case in a team meeting, is enough to show, later, that the safeguard was actually operating rather than existing only on paper as a policy nobody follows. Services sometimes build detailed alarm-fatigue policies specifying review intervals and escalation routes, and then produce no evidence, months later, that any of it actually happened for a given case; a light, consistent trace of the ordinary conversations described in this lesson is more useful than an unused formal procedure, because it shows the safeguard was alive in practice rather than only in the policy folder.

Sustained chronic risk work has a cumulative cost to the practitioners doing it, and guarding against drift is partly about protecting them, not only about protecting the people in their care. A service that treats alarm fatigue purely as a risk-management problem, without attending to the working conditions, caseload design, and support that make it more or less likely, is addressing a symptom while leaving its cause in place. Practitioners who raise early concern about their own capacity to keep holding a demanding chronic caseload are not reporting a failing; they are reporting exactly the kind of information a service needs in order to redistribute load, adjust supervision frequency, or bring in the structural safeguards described above before drift sets in rather than after it has already cost something. Treating that disclosure as a sign of poor resilience, rather than as useful operational information, teaches a whole team to stop disclosing it, which returns the service to relying on the one thing this lesson has argued cannot be relied

upon: an individual practitioner noticing their own habituation from the inside. The module that follows this one takes up the supervisor's and the organisation's responsibility for that side of the work directly.

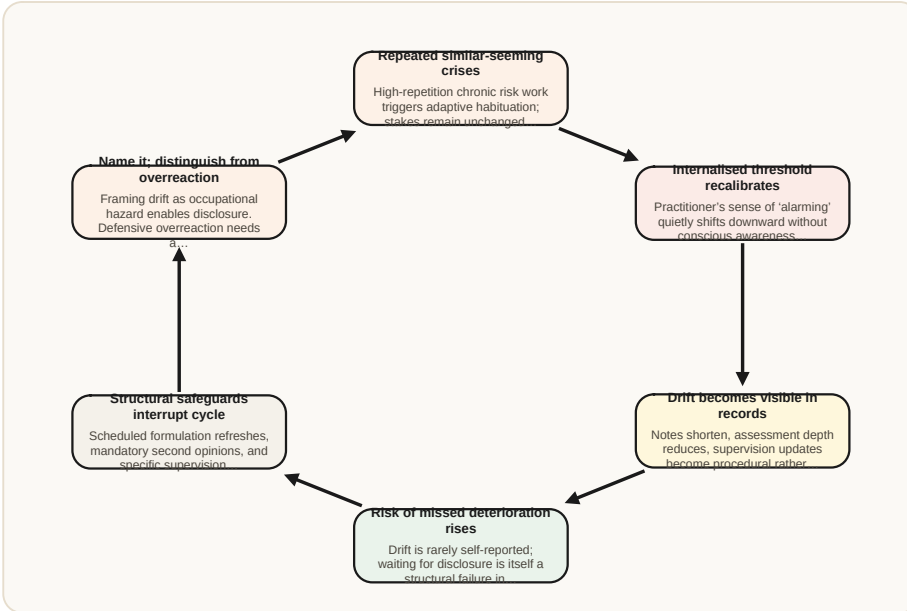


FIGURE 6.1 Managing Alarm Fatigue in Chronic Risk Work

Alarm fatigue is a predictable habituation cycle, not a character flaw. Structural safeguards interrupt the cycle at each stage before drift becomes a missed deterioration.

Points worth keeping

- Alarm fatigue is a predictable habituation effect of repeated, similar-seeming crises, not a character flaw, and naming it this way makes it something to manage rather than something to hide
- Watch for drift in note-taking, assessment depth, and supervision discussion becoming shorter or more procedural over time, rather than waiting for self-report
- Build structural safeguards that do not depend on a practitioner noticing their own drift: scheduled formulation refreshes, mandatory second-opinion reviews, and supervision that asks specific questions rather than seeking a general update
- Treat defensive overreaction as a related but distinct mechanism, often driven by vicarious anxiety rather than habituation, needing a different

supervisory response

PRACTICE TASK

Review your own recent notes on a long-standing case for signs of compression (shorthand replacing specific detail) across the last six entries. If you find it, identify one specific question you stopped asking that you could reinstate at the next contact.

The end of the sample

This is a free sample of “Advanced Risk Formulation and Safety Planning Practice”, and it ends here.

Everything you have read here is from the opening of the book, set exactly as it is printed in it. The rest follows in the same way, chapter by chapter, with its worksheets, its examples and its reference list.

The full book is on sale in the Talking Therapies UK shop. A copy is yours to keep, as a PDF and an ebook, licensed to you by name: a short line saying so is printed at the foot of each page of the PDF and on the title page of the ebook.