
Target panic as the de-automatization of a ballistic act

Matteo Campagna — MCPerformance · matteocampagna.com

Hypothesis first set out on 17 July 2026

ABOUT THESE PAGES

These pages present, in plain and extended form, a hypothesis about target panic: that it is the de-automatization of a ballistic act. It is a well-motivated model, built on established findings from the neuroscience of movement, but not yet directly verified.

A gesture that dissolves

In archery there is a failure that the vocabulary of sport cannot quite name, and that anyone who has not lived through it struggles to believe is real. An archer who has loosed tens of thousands of arrows reaches full draw: the anchor in its place, the sight steady on the gold, a single act left to perform, to open the fingers and let the string go. And the string does not go. Or it goes too soon, torn away before the aim is complete. Or it is the aim itself that no longer holds the center, and slips out of it every time it comes near.

From the outside it is a small thing: an arm that hesitates, an arrow gone crooked. For the person shooting it is another matter. The gesture that breaks down is not just any technique, but the one the athlete masters better than any other, repeated so many times that it has settled below the threshold of consciousness; and the moment they turn to look at it again, it crumbles in their hands.

Keep this detail in mind, because it is the thread that ties together everything that follows. We are used to thinking of mastery as something solid: the more you train a gesture, the more you possess it, the more you put it out of harm's way. Target panic says the opposite. The most automatic gestures are also the most fragile the moment attention returns to rest upon them: what makes a movement expert (the fact that it runs by itself, unsupervised) is exactly what exposes it to breaking as soon as supervision returns. At this level competence is not a wall built brick upon brick, but a balance; and a balance can be lost.

It does not concern a few unlucky people. The figures are uncertain, because it is a disorder that people confess only with difficulty, but in golf (where it goes by the name of *yips*, and it is the same thing) estimates of the symptoms vary widely from one study to another, from about one case in four to more than half among high-level players (Clarke, Sheffield and Akehurst, 2015); and they remain common among archers too. The same disorder appears in darts, in cricket, and outside sport in the cramp of those who write by hand and of musicians. The cost is proportionate. For an expert athlete the affected gesture is not one performance among others, but a settled part of their sporting identity: when it gives way, it can end a career and spoil forever the relationship with this sport.

Faced with such a phenomenon, research has split in two, along an ancient line: does the disorder arise in the body or in the mind? One tradition reads it as a dystonia, a fault in the movement in which muscles that should take turns contract together, and the limb stiffens against itself. The other reads it as an extreme form of *choking*, the collapse under pressure: anxiety brings conscious control back onto a gesture that has become automatic and breaks it into pieces, like the centipede that is asked in what order it moves its legs and that, by dint of thinking about it, stumbles. Each of the two brings concrete evidence. Each explains half of the phenomenon. And this is why, for decades, the field has been unable to decide.

The trouble is that neither of the two, taken alone, accounts for four traits that return in every description. The disorder appears all at once, without a decline to herald it: intact on one shot, in pieces on the next. It is tenacious: it lasts years, resists treatment, comes back in another form after an apparent recovery, and surfaces even in training, alone, with no audience and nothing at stake. It is of surgical precision: it strikes only the release and leaves everything else intact, draw, aim, hold. And it worsens with effort: the harder the athlete tries, and the more closely they watch the gesture, the worse it goes. The nerve explanation stumbles on the abrupt onset and on the appearance in solitude; the muscle explanation stumbles on the fact that the block dissolves when shooting into a target with no center and returns the moment a score to defend comes back. Each holds one half in hand, and the other refutes it.

In these pages I argue that those four oddities are not separate riddles, but the signature of a single mechanism, and that to recognize it one must stop looking for the broken part. My proposal is that nothing breaks: not the muscle, not the nerve, not the motor program. The release that fails is the same one that succeeds, executed however in the wrong mode of control. I call it *de-automatization*: under attention and under the threat of a bad outcome, a rapid gesture, which the body performed by predicting it and letting it go, regresses to the vigilant and rigid control that the motor system adopts every time something is uncertain or costly to get wrong. It is a mode that a gesture this brief has no time to use, and from which it can only come out ruined. Read this way, the psychological origin and the physical signature are not two causes in competition, but cause and effect of a single process: the estimation of threat, and the stiffening it summons onto a gesture that cannot afford it.

I want to say at once what all this is and what it is not. It is a hypothesis. It rests on solid findings from the neuroscience of movement and holds up well against the anomalies of the disorder, but agreement with the facts is not proof: around the same data one can build many elegant stories, some of which may turn out to be false. The value of this one does not lie in the breadth of what it manages to explain, explaining much comes cheap, but in the fact that it can be put to the test, and, if it is wrong, shown to be so. For this reason, after building it, I devote a whole part to exposing its flank: to listing the precise findings that would sink it, and the experiments, almost all feasible with instruments already in use, capable of producing them.

The path is this. In the first part I look closely at the existing explanations (dystonia, choking, the continuum that tries to add them up, the catastrophe model, the ecological reading) and I measure their common gap. In the second I put on the table the few facts about how the release is really controlled: how fast it is, why it cannot be corrected by watching it, why the body performs it in an almost automatic manner. In the third I state the hypothesis and lead it through the four anomalies. In the fourth I expose it to the fire of evidence: seven predictions, each built so that it can fail. In the

fifth I say where it may be wrong. In the last I return to the archer at full draw, to say what changes, if the hypothesis is right, for those who shoot and for those who teach.

Part I • The gesture and the explanations that are not enough

Research on the yips revolves around a single disagreement, and the disagreement is deep. Before proposing a way out it is worth looking carefully at the explanations that already exist: not to discard them, but because the hypothesis to come makes them fit together. There are five, and each grasps something true. The problem is what each leaves out.

1 • The body: focal dystonia

The first tradition places the disorder entirely in the body. The yips, it holds, is a focal dystonia: a fault in the movement related to writer's cramp and to the musician's cramp, those disorders in which a hand that has repeated the same gesture for a lifetime at a certain point contracts by itself. In a dystonia the muscles that should take turns fire together. This is *co-contraction*: the agonist, which carries the limb in one direction, and the antagonist, which brings it back, pull at the same instant; the limb, contested by two opposing forces, locks or jolts against itself. It is the same principle as the accelerator and the brake pressed together: the engine pushes, the brakes hold, the car stays still and vibrates. In the release, that contest between antagonist muscles is what the archer feels as block, snatch, tremor.

It is not a conjecture pulled from thin air. It arose, at the end of the 1980s, from a clinical observation on golf: the affected players turned out to be older and more expert than the healthy ones, and in a quarter of cases the disorder extended beyond golf, the profile of an occupational dystonia, the price that certain hands pay for having repeated the same fine movement too long. A few years later direct evidence arrived: recording the electrical activity of the forearm muscles during the putt, anomalous co-contraction was found in half of the affected golfers, and in none of the healthy ones taken as controls. The signature was there, and it was only in the affected.

This reading explains two things well. The physical face of the disorder, first of all: the snatches, the blocks, the tremors are not the imagination of a tense athlete, but a limb that really stiffens, and now we know by what mechanism. And its specificity: an occupational dystonia strikes a single over-trained gesture and spares everything else, exactly like target panic.

The problem is another, and it is the most serious one for this theory. If the disorder were a fixed lesion in the way those muscles are commanded, it should manifest every time the gesture is performed. Instead it comes and goes with the situation. Put that archer in front of a bale of straw with no center, with nothing to hit and nothing to get wrong, and often the block dissolves: the string goes clean. Put back a target with scores, or even just someone watching, and the block returns. A lesion in the muscle does not switch on and off with the arrival of a spectator. Whatever target panic is, it is the situation that commands it, and this, on its own, is enough to say that it cannot live in the muscle alone.

2 • The mind: choking and reinvestment

The second tradition makes the opposite argument: the disorder does not arise in the muscle, but in attention. It is the chronic and extreme form of *choking*, that collapse under pressure that every athlete knows at least by hearsay. Its most elaborate version is called *reinvestment theory*, and its central idea is as simple as it is convincing.

A well-learned skill runs by itself, without the athlete following it step by step. The gesture has been repeated so many times that it has become a closed procedure, which runs beneath consciousness and no longer needs explicit instructions. Then pressure arrives, and with pressure attention returns to the gesture. The athlete begins again to control it consciously (to put their head into it, to tell themselves what to do while doing it) and in doing so breaks a continuous routine down into the hesitant control, one piece at a time, of a beginner. It is the same movement, but governed by the wrong part of the system: working memory, which reworks explicit rules, takes the place of the automatic procedure that was working.

Try going down a flight of stairs while describing aloud, step by step, how you shift your weight: you slow down, you hesitate, and for a moment you walk worse than a child. Or think about how you type on a keyboard: your fingers know where the keys are until you try to recall, one by one, where each letter is. Explicit attention does not add control to an automatic skill: it dismantles it.

This reading explains well the disorder's sensitivity to pressure, and above all it grasps the paradox at the heart of everything: it is precisely watching the gesture that ruins it. I will return to this, because it is the right intuition, the one on which the hypothesis is built. What this tradition does not explain is the rest. Not the sudden onset, because anxiety usually mounts by degrees, not in jumps. Not the appearance of the disorder in a quiet training session, alone, when there is no pressure anywhere. And not its openly motor face (the block, the snatch, the tremor) which resembles a dystonia far more than a mere state of mind.

3 • The continuum: a map, not an engine

With two theories each grasping half of the phenomenon, the field did the thing one usually does in these cases: it lined them up along a *continuum*. The yips would not be a single disorder, but a spectrum, with a predominantly neurological pole (dystonia) at one extreme, a predominantly psychological pole (choking) at the other, and in between the mixed cases, which show a bit of one and a bit of the other. The same idea was then arranged into a typology: Type I manifestations, physical; Type II, psychological; Type III, both.

As a clinical tool, for matching an intervention to the type of case in front of you, all this is useful, and remains useful. But as an answer to the question of what the yips is, a continuum is a map, not an engine. It tells you where to place a case between two extremes; it does not tell you how the disorder is produced. And there is worse: by distributing the cases into physical, psychological and mixed, it takes for granted precisely the separation between mind and body that it should serve to overcome. The initial disagreement has not been dissolved. It has been filed away in three boxes.

4 • The catastrophe: the abrupt collapse and the hysteresis

There is a third line of research that usually does not appear in discussions of the yips, but that here must be looked at, because it establishes a fact we will need: that two of the traits we are trying to explain (the abrupt and discontinuous collapse, and its resistance to going back) are real properties of expert performance under pressure, not the impressions of shaken athletes.

It arises from an observation. Athletes who go beyond their limit do not worsen gradually: they collapse all at once, with a wide and sudden drop, and then struggle to climb back. To account for this a *catastrophe model* was proposed, borrowed from mathematics: performance as a discontinuous function of two variables, cognitive anxiety (worry, thoughts about the result) and physiological arousal, the body's activation. As long as cognitive anxiety stays low, arousal and performance follow the old inverted-U curve: one performs better at an intermediate level of activation, worse when too relaxed or too keyed up, and the transition is gradual. But when cognitive anxiety is high, the curve folds: performance rises with arousal up to a certain point and then plummets, all at once, as if off a cliff.

The signature of this model is *hysteresis*, and it has been confirmed experimentally. Hysteresis means that the way down is not the same as the way up. Under high cognitive anxiety, once performance has plummeted, it is not enough to bring arousal back to the level at which it fell to recover it: arousal must drop much lower, well below that point, before performance picks up again. It is the asymmetry shown in the figure, and it does not appear when cognitive anxiety is low.




Figure 1. The catastrophe collapse. With low cognitive anxiety, performance and arousal follow a soft inverted U. With high cognitive anxiety the curve folds: performance rises, then plummets all at once, and on the way back arousal must drop well below the collapse point before it recovers, the signature of hysteresis. Original schematic representation for teaching purposes. The curves illustrate the qualitative behavior of the model, not experimental data.

The catastrophe model thus demonstrates, with evidence in hand, that performance under pressure can change all at once and behave in a hysteretic manner. But it does so at a level that, for our problem, is the wrong one. It describes overall performance as a function of anxiety and arousal; it is not a model of the motor mechanism of any particular gesture, and it does not specifically concern the

yips. It tells us that the dynamics we are looking for (the abruptness, the hysteresis) are genuine, but it does not locate them within the control of movement. It establishes that the phenomenon is real without saying where, in the gesture, it happens. It is a piece we will keep, and to which we will give a mechanism further on.

5 • The ecology of skill: the disorder that lives in the context

A more recent line shifts the gaze again, and in a direction different from the first two. Instead of looking for the disorder inside the athlete, in the muscle or in the mind, it reads it as something that happens between the athlete and the world. It arises from a study of archery conducted from the inside, by someone who shoots, and describes target panic as the disintegration of what it calls the *ecology of skill*: the web of relationships between the archer, the equipment, the target and the environment through which an expert gesture normally holds itself together. The warning it carries is precise, and it is aimed against an ancient habit: the Cartesian one of looking for the cause of the disorder only inside the person, as if it were a fault localized in an organ or in a state of mind.

The merit of this reading is that it puts at the center the trait that most embarrasses the muscle theory: the dependence on context. It is the fact, already encountered, that the block dissolves when shooting into a neutral target, with no center to hit, and returns the moment a score or a gaze comes back. However one wishes to read the disorder, this observation on its own establishes something: that it is the situation that switches it on and off, and that therefore it does not simply reside in the muscle.

It is a true and important observation, and I will keep it. But it points to a problem more than it solves one. To say that context switches the disorder on and off is not yet to say which mechanism, inside the archer, context is switching on and off. The dependence on context is the fact to be explained; it is not yet the explanation.

6 • The gap at the center

Put together, these explanations leave a hole with sharp outlines. The two poles each explain half of the phenomenon and are refuted by the other half. The continuum orders the cases without saying how they arise. The catastrophe model demonstrates that the abrupt and hysteretic collapse is real, but describes it at the wrong level to explain why it strikes one gesture and spares another. And the ecological reading establishes that context switches the disorder on and off, without saying which mechanism context is commanding. The table summarizes the position.

Tradition	What it explains well	What it leaves out
Focal dystonia · the body	The physical signature (snatches, blocks, tremors) and the specificity to a single gesture; co-contraction measured in the affe	cted. The dependence on context: why the disorder disappears on the neutral target and returns with the score.
Choking and reinvestment · the mind	The sensitivity to pressure and the central paradox: watching the gesture breaks it.	The abrupt onset, the appearance in solitude, the motor signature resembling dystonia.
The continuum · Type I / II / III	A useful clinical guide for matching intervention and case.	It is a map, not an engine: it places a case, it does not explain how it arises.
The catastrophe	That the collapse is really abrupt and hysteretic, not an impression.	It describes overall performance, not the mechanism of a single gesture.
The ecology of skill	That it is the situation that switches the disorder on and off.	It does not say which mechanism, inside the archer, the situation commands.

What is missing is a single hypothesis, placed at the level of motor control, that does three things at once. That produces from a single mechanism both the psychological trigger and the physical signature, instead of keeping them as two separate causes to be reconciled. That explains the four anomalies (the abrupt onset, the persistence, the specificity, the worsening with effort) from the same principle. And that locates, within the control of movement, that sudden, tenacious and context-bound collapse that the other theories either do not explain or describe from too far away.

Building this hypothesis is the task of the rest of these pages. And to build it, we need first some facts about how the gesture that breaks down is really made.

IN BRIEF

Five explanations, each with a part of the truth. The body (dystonia, measured co-contraction) and the mind (choking, reinvestment) each grasp half of the phenomenon and contradict each other; the continuum adds them up without fusing them; the catastrophe model proves that the collapse is really abrupt and hysteretic, but from too high up; the ecological reading shows that it is context that commands, without saying what. What is missing is a single mechanism, at the level of motor control, that produces trigger and signature together and dissolves the four anomalies.

Part II · How the gesture is really made

To find where the release breaks, one must know how it works when it does not break. Not much is needed: three facts suffice, all established in movement research. The first is the difference between two ways of controlling a gesture, and the time constraint that decides which of the two a given gesture can use. The second is that the expert control of a rapid gesture is not only predictive, but also soft. The third is that the same system, every time it judges a movement uncertain or costly to get

wrong, does the opposite: it stiffens. These are the three pieces that, put together, open the door to the disorder.

7 • Two ways to control a movement

There are two ways to hold a gesture steady, and the difference between them is the key to everything. In the first (the *predictive* control, which engineers call open-loop) the command for the movement is computed beforehand, starting from an internal model of the limb and the target, and then launched without waiting for the senses to say how it went. In the second (the *feedback* control, closed-loop) the movement is monitored while it happens, continuously compared with what vision and proprioception (the sense of the body's position) report, and corrected on the fly.

It is not a clean choice between the two. In today's most accredited model of how the nervous system governs movements, it does not impose a rigid trajectory to follow, but decides instant by instant how much weight to give the sensory correction: it raises its listening to feedback where correcting serves the target, it lowers it where it does not. And a gesture aimed at a target uses both modes, in sequence. A first phase, already programmed, launches the limb toward the center; then a feedback phase guides the final approach, the fine-tuning.

It is exactly the structure of an archer's shot. The long aim and the hold (the sight kept on the gold, the small continuous adjustments) are full of feedback: vision and proprioception inform ceaselessly, and the archer corrects. The release is another matter. And to understand why, one must look at the clock.




Figure 2. The two ways to control a gesture. On the left, predictive control: the command is computed beforehand and launched, with no time to correct. On the right, feedback control: the movement is monitored and corrected while it unfolds. Aim and hold use them in sequence; the release, as will be seen, does not. Original diagram for teaching purposes.

8 • Feedback is too slow

Feedback has a price, and it is time. A signal that leaves the senses must be collected, transmitted to the central nervous system, processed and turned into a correction command; and this whole loop is

not instantaneous. The loop that carries visual information about the limb to the point of being able to modify a movement already under way has a delay on the order of one hundred, one hundred and fifty milliseconds. It is a very short time by our mental clocks, but it is not zero, and for certain gestures it makes all the difference.

The consequence is stark. Any movement that concludes in less than that time cannot be guided by vision while it happens, because the loop has no time to close: by the time the information comes back with the correction, the gesture is already over. It is like driving while looking only at the rear-view mirror, which shows the road of an instant ago: fine for correcting a long curve, useless for dodging something that appears in front of you in a fraction of a second.

And it is precisely to solve this problem, how the brain manages to govern movements faster than its own feedback, that the nervous system relies on *predictive internal models*. An internal model is a simulation: it predicts the consequences of a command before the senses confirm them, and takes the place of the feedback that would arrive too late anyway. The brain, for fast gestures, does not guide by watching: it guides by predicting.




Figure 3. The window of time. The visual feedback loop takes about 100 to 150 ms to close; the release lasts a few milliseconds. The gesture is over long before vision can correct it: it cannot be piloted, it can only be launched.

Original schematic representation for teaching purposes. Values representative of the trends, not original datasets.

9 • The release is ballistic

The release is precisely one of those gestures. The aim and the hold that precede it can be corrected with feedback, and they continuously are. But the release itself (the fingers opening, the string going) is a brief and *ballistic* event: launched and concluded in a flash. The finger joints rotate and free the string in a fraction of a second, too quickly for the gaze to follow and correct it.

The whole event is much briefer than the feedback loop. There simply is not the time to watch the release and correct it while it happens. And this does not depend on the technique of this or that archer: it is a necessity of time. The gesture is faster than the loop that would be needed to monitor it, and therefore it must be executed in predictive mode, pre-programmed and launched, not piloted.

I want to be precise, because it is a point on which one can be misunderstood. I am not saying that no feedback reaches the release: real movements always mix predictive control and fast internal, proprioceptive returns, and the release is no exception. The claim is narrower: it is *visual and real-*

time correction that time forbids, not every form of feedback. Whatever else is true of the release, it must start already decided.

And it is also the reason why the old teaching tradition repeats that the release is not made, but let happen. It is not a poetic turn of phrase: it has a precise basis. Every attempt to deliberately control, instant by instant, the departure of the string is the attempt to do something that time does not allow. The masters, here, were right without fully knowing why.

10 • The expert gesture is soft; the system stiffens under threat

Time is not the only thing that matters. A gesture can be held steady in two ways. One is to correct it while it happens, the way that at the release is forbidden. The other is to stiffen the limb that produces it, contracting the opposing muscles together to increase its *mechanical stiffness*: a stiff limb resists disturbances by itself, without needing to perceive and correct them. It is the strategy you use, without noticing, when you carry a full glass in your hand over uneven ground: you stiffen the arm, and the glass stays still.

The decisive point is when the motor system chooses this way. It chooses it when prediction is not yet enough. Faced with a new or unstable movement, it first stiffens, raising co-contraction to oppose by brute force a disturbance it cannot yet predict; then, as it learns an accurate internal model, it lets the stiffening drop and releases the limb. Co-contraction is therefore the mark of a movement *not yet, or no longer, automatic*: high when the skill is new, low when it has consolidated and the internal model has taken command, restoring to the limb its softness. The expert release is, by this measure, the soft one.




Figure 4. Alternating activation versus co-contraction. On the left, the normal pattern: agonist and antagonist take turns, the limb moves freely. On the right, co-contraction: the two muscles fire together, and the limb stiffens against itself. It is the signature measured in affected golfers. Original illustration for teaching purposes.

Two facts complete the picture, and they are the ones the hypothesis will lean on. The first: how much the system stiffens does not depend only on the mechanics of the gesture, but on its estimate of how much could go wrong. Stiffness rises with uncertainty and with the cost of an error, and, crucially, psychological stress alone is enough to raise it. No physical load is needed: the perceived stakes suffice to make the arm stiffen. The second, which will serve only for the form of the onset and not for the mechanism: the motor system is capable of truly abrupt, threshold-like transitions between one coordination pattern and another, as has been seen in studying the rhythmic movements of the two hands. The jump, the sharp change, is already in its repertoire; there is no need to invent it.




Figure 5. Stiffening across the arc of learning. Co-contraction is high when the gesture is new, falls as the skill consolidates until an expert release, soft and automatic, and can return, under attention and threat. It is this return that I call de-automatization. Original schematic representation for teaching purposes. Values representative of the trends, not original datasets.

11 • The pieces on the table

Three facts are now on the table, and it is worth aligning them, because together they open a door. The release is too rapid to be corrected by feedback in real time, and therefore it must be executed by predicting it. The expert execution of a similar gesture is also soft: the relaxed control, based on the internal model, that practice installs. And the motor system carries with it a permanent disposition to do the opposite, to switch monitoring back on and to stiffen the limb, every time it judges high the stakes or the uncertainty of a movement.

Placed side by side, these three facts make available a fault of a particular kind. The control that the release requires (predictive, soft) and the control that attention and threat summon (monitored, rigid) are opposites. And the second can be applied to the release even if the release cannot use it. Read this way, the disorder would be neither a fault in the muscle nor a collapse in the nerves, but the expert control of the gesture displaced by a mode of control right for slow or uncertain movements and ruinous for this one. It is this possibility that the following part turns into a hypothesis.

IN BRIEF

Three facts. First: the release lasts a few milliseconds, far less than the 100 to 150 ms of the feedback loop, so it cannot be corrected on the fly, it must be launched by predicting it. Second: the expert control of a similar gesture is also soft, because co-contraction falls as the internal model installs itself. Third: the same system raises monitoring and stiffness every time it estimates high the cost of an error, and psychological stress alone is enough to do it. The control that the release requires and the one that threat summons are opposites.

Part III • The hypothesis: de-automatization

12 • The thesis

Now I can state it in full. **Target panic is the de-automatization of the release:** under attention and under the threat of a bad outcome, the gesture regresses from the predictive and soft control that skill had installed to the monitored and rigid control that the motor system applies when something is uncertain, a mode of control that a gesture this brief is too fast to use, and from which it can only come out ruined.

The move at the center of everything is simple, and it is worth isolating. The disorder is not a broken part, but a mode of control healthy in general and wrong for this gesture. Nothing needs to break: not the muscle, not the motor program, not the nerve. The release that fails is the same release that succeeds; what has changed is the way it is executed. The system has regressed from predict-and-let-go to monitor-and-stiffen, and the gesture cannot bear being either monitored or stiffened without breaking down.

From here the forms of the disorder read by themselves. The block is the limb stiffened against its own movement. The snatch and the disorganized release are what remains of a ballistic gesture when a late and mistimed control is injected into it. And the co-contraction that the dystonic tradition measures is not a lesion to be reconciled with difficulty to psychology: it is exactly what psychology predicts, the stiffening that a motor system produces when it treats a movement as dangerous.

In explicit form, the proposal rests on four claims. The first three have already been established; the fourth is the new step.

- **The release must be executed predictive and soft.** By the necessity of time it is too rapid for on-the-fly correction, and the expert execution of a similar gesture is the relaxed control, based on the internal model, that practice installs.
- **The system has a disposition to stiffen.** Permanent and usually useful: it switches monitoring back on and raises co-contraction every time it estimates high the cost of an error.
- **Attention and threat are an estimate of that kind.** The impulse to watch the string go and to control its departure, and the weight of the score, summon precisely that monitoring and that stiffening onto the gesture.
- **Applied to the release, they degrade it.** Monitoring arrives too late to be of use, and stiffening opposes precisely the limb that must move freely. The gesture, executed in a mode it cannot use, breaks down in the attempt. The disorder is this misapplied mode, de-automatization, and its motor signature is co-contraction.

13 • The four anomalies, dissolved

A single mechanism, and the four oddities of the first part stop being separate riddles. Each falls out of a precise trait.

Sudden onset. The onset can be abrupt without needing a jump between two states. Co-contraction rises continuously with perceived threat, but its effect on the gesture is not continuous: below a

certain threshold the release still escapes the stiffening and goes clean; above it, the antagonist blocks the joint and the gesture freezes or is torn away. A gradual change in the hidden variable shows up as a step change in the visible gesture. And it can be even sharper if the association between a cue and threat forms in a single occasion, the bad shot that the athlete indicates as the moment when it began, because single-trial aversive learning can install a conditioned response in one go.

Persistence. The heightened control does not subside when the pressure of the moment passes, because what sustains it is not present anxiety, but a learned association. The cues of the shot (the center, the score, the felt moment of the hold) come to summon the monitor-and-stiffen response on their own. This is why the disorder lasts between one session and the next, resists the invitation to relax, and appears even in a calm and solitary training session: the trigger is a conditioned cue, not the state of nerves of the moment.

Surgical specificity. Only a gesture too rapid for on-the-fly correction, and too dependent on a free limb to tolerate stiffening, is damaged by the regression. The surrounding skill (the draw, the aim, the hold) is slow enough to use monitoring and robust enough to withstand a bit of stiffness, and for this reason it is spared. The specificity is not an extra hypothesis: it descends from the requirements of time and mechanics of the release.

Worsening with effort. The trigger is the perceived need to control the gesture. Trying harder means paying more attention and desiring the outcome more: both raise the estimate of risk, and with it the monitoring and stiffening that ruin the gesture. The relationship is graded and self-harming by construction, more effort, more control, more breakdown. It is the paradox that every theory has noted and none has explained; here it is not a curiosity, but a direct consequence of what the trigger is.

Anomaly	How it presents	Explanation in the model
Sudden onset	Intact on one shot, in pieces on the next.	A threshold on a variable that rises continuously: below, the release goes clean; above, it blocks. Or an association learned in a single bad shot.
Persistence	Lasts years, resists the invitation to relax, appears even alone and with no stakes	. The trigger is not present anxiety, but a conditioned cue (center, score, moment of the hold) that summons the stiffening on its own.
Specificity	Strikes the release, leaves draw, aim, hold intact.	Only a gesture too rapid for correction and dependent on a free limb dies in that mode; the rest is slow and robust enough to survive.
Worsens with effort	The harder you try, the worse it goes.	More effort = more attention + more desire for the outcome = higher estimate of risk = more stiffening. Self-harming by construction.

Table B. The four anomalies, and how a single mechanism, the estimation of threat and the stiffening it summons, produces them all.

14 · The two traditions as cause and effect

At this point the reconciliation between body and mind is no longer a compromise, but a consequence. The two traditions were measuring two parts of the same process. The psychological tradition identified the *trigger*: anxiety, pressure and reinvestment are all paths by which attention and the cost of error turn onto the automatic gesture. The dystonic tradition measured the *motor consequence*: the co-contraction recorded in the affected. Not two causes in competition, but cause and effect within a single mechanism, the heightened estimation of threat, and the stiffening it produces.

That co-contraction is not the imprint of a lesion: it is the same stiffening response that a healthy motor system enacts under uncertainty everywhere, here summoned onto a gesture that cannot afford it. This is why each field found true evidence and neither could reduce the other to itself. And this is why the dependence on context, the disorder that dissolves on the neutral target and returns when a score or a gaze comes back, is exactly what the hypothesis predicts: context is what raises or lowers the estimation of threat.

On one point it is worth being explicit, because the two readings risk confusing it: co-contraction plays two different roles here. It is the observable signature of the de-automatized mode (the motor sign, read by electromyography, of a system stiffened by threat) and at the same time I advance it as the proximal causal mediator of the fault: not only what is seen when the release fails, but that through which the failure passes. The two things are different (one is an observation, the other a cause) and must not be confused: that co-contraction accompanies the disorder would be said by the dystonic thesis too, whereas that it is the mediator is the stronger and riskier claim, the one that only manipulation can decide. And it is exactly what Prediction 7 puts to the test: imposing stiffening without threat should make the deficit appear, removing it under threat should save the release.

It is worth saying, and not hiding among the limits, a question that the mechanism leaves open: the wrong mode ruins the release by two distinct paths. One is the limb stiffened by co-contraction, which brakes the displacement and, at the limit, blocks it: it is the path of the block and of hypometria. The other is the late corrective commands, injected by feedback into a gesture that is already over: it is the path of the snatch. Which of the two carries most of the failure, and how they combine, I do not decide a priori: it is a matter to be measured. It is settled by two converging paths (the dissociations of Prediction 3 and Prediction 7, where the first isolates the corrective path, tied to latency, and the second the impedance path, the mediator) and by putting the mechanism to the test in a model that switches one channel on and off at a time. Such a model, with parameters taken from the literature and not tuned on the result, suggests that the two channels divide the work: stiffening produces the graded hypometria, while the specificity to duration is carried by the corrective path, which repairs the deficit only when the gesture lasts longer than the loop's latency.

The hypothesis also explains a clinical fact that has frustrated both traditions: that lowering arousal (the body's activation, with relaxation or with beta-blockers) so often does not relieve the disorder. The variable that commands is not arousal, but the cue-driven estimation of threat and the mode of control it summons; calming the body does not necessarily lower that estimate, nor does it free a limb that a conditioned cue is still stiffening. Not by chance, the evidence that psychological stress raises limb stiffness comes precisely from the reinvestment tradition: stiffening is the bodily form that

reinvestment takes. And the interventions that do help (a neutral target, an external focus, a modified routine) are those that lower the estimation of threat or shift control away from the gesture.

15 • What it adds to reinvestment

It is worth being precise about what this hypothesis adds to the reinvestment theory on which it rests, and about what instead it owes to it, because the two are easily confused, and because not everything here is mine. I say it openly: the idea that the yips is a de-automatization does not originate in these pages. The term itself is older, borrowed from a 1960s work on attention, and it has long been the assumption on which reinvestment theory rests; and the thesis that it is precisely reinvestment, the conscious control that displaces the automatic skill, that causes the disorder has already been advanced, and examined directly on the yips. On this I claim nothing. Reinvestment, and alongside it explicit monitoring theory, identify the trigger and its general effect: conscious control is reinvested in an automatic skill, and the skill breaks down. It is the starting point, and it is right. What those theories do not specify is the *motor mechanism* of the breakdown, and with it three things that remain unexplained.

Why precisely this gesture, the release and not the draw? Because only an act too rapid for on-the-fly correction, and too dependent on a free limb, is ruined by the monitoring and stiffening that reinvestment summons. Why so suddenly? Because a stiffness that rises continuously crosses a threshold that blocks the gesture, and because the association that triggers it can be learned in a single occasion. And why does it stay broken, able to survive the pressure that set it off? Because the heightened control conditions itself to the cues of the shot. Reinvestment tells us *that* attention breaks the skill; this hypothesis tells us, in the language of motor control, *how*, and therefore why this skill, why all at once, why it does not heal on its own.

16 • The catastrophe, explained from within

Let us return, with one more piece, to the catastrophe model. The sudden and hysteretic collapse of performance under high anxiety, which that model proposed and the experiment confirmed, is, in this reading, the shadow at the level of overall performance of the de-automatization of specific gestures at the level of motor control. The catastrophe model had grasped the *form* of the failure: its abruptness, its resistance to a simple lowering of arousal. But it had grasped it from afar, without a mechanism. This hypothesis supplies the mechanism that has that form and locates it in the control of the gesture. What was a description becomes, if the hypothesis is correct, an explanation.

17 • The chronic course: the same control, conditioned and avoided

The hypothesis, so far, concerns acute failure: the mode of control summoned onto a single shot. The course that the disorder follows over months and years is the same mechanism extended in time, and it follows from a fact already in play: the heightened control is summoned by a cue, and cues are learned.

Target panic is described as chronic and resistant to treatment, and beyond simple persistence it tends to deepen: the point at which the shot is lost migrates ever earlier in the sequence (from an inability to hold the aim on the center, to an inability to hold at all, to a release that escapes even before the anchor) until, in advanced cases, it appears with no target at all or with the eyes closed.

Two learning processes, on the same control parameter, produce this course. The first I give with a certain confidence; the second I flag as more conjectural.

The first is a simple conditioning of the mode of control. If the monitor-and-stiffen response is evoked continuously in the presence of the center, the score and the felt moment of the hold, those cues come to evoke it on their own, and to do so ever earlier and in ever more numerous contexts as the association widens, until the response appears with no center to aim at and no score to fear. It is an established reading: the yips as a context-bound movement disorder, developing in a vicious circle of anxiety, over-control, interference and awareness in which each factor feeds the others, and in which the movement is often normal in a sufficiently different context, the signature of a learned, cue-bound response, not of a fixed lesion.

The second process, more speculative, is an *avoidance* dynamic driven by negative reinforcement. Once the moment of the hold has become unpleasant, discharging the shot early removes the athlete from that moment; the escape brings relief, and the relief reinforces the escape, so the release goes even earlier, shot after shot. It is the natural reading of the snatch and of the backward migration, but it remains a proposal more than a datum, and I keep it separate from the rest.

The two readings, acute and chronic, point to the same training: do not add control to the release, and do not let the athlete keep escaping from it. Both amount to rebuilding the conditions in which the release can again be let happen, instead of controlled until it happens. And what distinguishes all this from a second theory in its own right is no small thing: here there are not two disorders, but a single control parameter, the level of monitoring and of stiffness on the gesture, raised all at once by attention and rooted over time by learning.

A hypothesis of this kind earns the right to be taken seriously only by exposing itself to refutation. The mechanism I have described is not seen directly; its consequences are seen, and a good hypothesis must point to consequences that would not follow if it were false. It is the task of the part that comes now.

IN BRIEF

A single mechanism, the estimation of threat and the stiffening it summons, produces the four anomalies, reconciles body and mind as cause and effect, and absorbs the catastrophe model by supplying its mechanism. Its chronic arm is learning: the mode of control conditions itself to the cues of the shot (solid) and, perhaps, the anticipated escape reinforces itself (more conjectural). A single control parameter, raised all at once by attention and rooted over time.

Part IV • The evidence: seven predictions, and one experiment

18 • A hypothesis is worth the risks it runs

A hypothesis is worth the risks it runs. The de-automatization of the release is not observed directly, but it has consequences that the hypotheses it means to replace do not share, and it is on those

consequences that everything is played out. I set out seven. Each is formulated so that a certain result would follow if the hypothesis is correct, and would not follow if the disorder were anxiety alone, dystonia alone, or a catastrophe at the level of performance without a localized mechanism. Where a prediction can be put to the test by adapting a method already in use, I say so; where a new one is needed, I describe its design.

In all of them, the criterion that discriminates is the same. My theory ties the disorder to the level of monitoring and of stiffness applied to a specific gesture, briefer than the latency of feedback, and predicts that it is the manipulations of that control, not of arousal, that switch the disorder on and off, and that do so in a graded manner. Since every prediction revolves around a single construct (threat, that is the expected cost of the error that the system estimates from the gesture) it is worth making it operational first, and then stating the predictions in its terms. It is the most technical passage of these pages, and it is also the one that holds up all the rest: without it, the central thesis risks going in circles.

19 • Defining threat without going in circles

The hypothesis revolves around a single variable, the threat that the system estimates from the gesture, and that variable must be made operational, or the central thesis collapses into a tautology. If threat is defined as *what worsens the release*, then saying that “threat raises the stiffening that worsens the release” is true by construction, and explains nothing. I therefore fix the construct in a way independent of the outcome. By threat I mean the estimate, by the controller, of the *expected cost of the error*: the perceived stakes of the gesture multiplied by the perceived probability of missing, the same quantity that the framework of optimal control and impedance already handles, because it is uncertainty and the cost of error that fix how much weight to give the feedback and how much to stiffen the limb.

And it is not a private feeling to be guessed after the fact. The expected cost of the error is a quantity that can be *imposed and measured before* the gesture: when it is made operational as an explicit structure of penalty and reward, it demonstrably reshapes the planning of fast, goal-directed movements, shifting the point at which one aims so as to maximize the expected gain. It is this variable, manipulable and measurable, that the predictions put to work.

The construct and its consequences occupy four distinct levels, and confusing them is precisely what generates the circularity.

- **(i) The inputs** fix the expected cost: the stakes tied to the shot, the salience of the score, the evaluative pressure, the attentional instructions, the probability of error set by the size and distance of the target. They are imposed by the experimenter, and independent of how the shot goes.
- **(ii) The antecedent measures** verify that the estimate has actually risen: validated self-reports of the perceived stakes and of cognitive anxiety, distinct from somatic anxiety, together with autonomic indices of arousal, such as heart rate variability, electrodermal activity and pupillometry, all recorded *before* the loose.
- **(iii) The mediators** are the motor and attentional signature that the theory predicts that estimate must produce: antagonist co-contraction (surface electromyography), limb impedance

(the response to a mechanical perturbation), the *quiet eye* (the last fixation on the target before the release, whose earlier onset or shift toward the string hand marks vision re-hooked to the gesture) and left-temporal to frontal coherence.

- **(iv) The outcome** is what degrades: the variability of the release timing, freezing and premature release, reaction time to the clicker, the score.

From here descends, straight, the *cardinal rule*: the construct is measured only at levels (i) and (ii), and never read from the mediator or the outcome. Defining threat through the co-contraction or the timing variability that it should explain would set the tautology back on its feet; those are consequences to be predicted, not measures of the antecedent.

Once the construct is fixed upstream, four *dissociations* become available that a circular variable could never produce. The first is the double dissociation by gesture duration: the same threat, imposed and measured, applied to the sub-latency release and to a slowed, correctable version of the same hand action, should degrade only the rapid gesture, with the antecedent indices equal in the two (this is Prediction 3, which thus becomes a test of non-circularity more than a curiosity). The second is arousal without appraisal: raising the body's activation without raising the perceived cost of error (with physical exercise, or with a stressor unrelated to the shot) should leave the release intact, separating threat from mere activation.

The third is the manipulation of the mediator: imposing stiffening without threat should reproduce the deficit, and reducing co-contraction under threat should save the release, so that the outcome follows the measured mediator and not threat, once the mediator is fixed (this is Prediction 7). The fourth is the most instructive: high measured threat with intact release. These are the conditions of high antecedent index but clean execution (the slow gesture, the expert who controls the release without losing it, the case of arousal alone) that the theory *requires* to exist and that a circular variable would forbid. The short answer to the objection that we would be calling threat everything that worsens the shot is therefore this: threat is defined by the imposed stakes and instruction, verified by appraisal and autonomic indices recorded before the loose, and the proof that the definition does not go in circles is that the hypothesis predicts, and can be brought to show, high measured threat with intact outcome. The design, in all this, is within-subject, within-session and reversible: switching the input on and off should switch on and off together the antecedent, the mediator and the outcome, in the same archer and in the same minutes, a coupling in time that is already in itself a proof against circularity.

20 • Prediction 1 • Co-contraction grows with the stakes

If the disorder is the motor system stiffening a gesture that it should leave free, the stiffening should scale with the expected cost of the error. As that estimate is raised (by making the stakes and the score more salient, or with instructions that direct attention to the hand and to the release) antagonist co-contraction during the loose should rise with it: graded, monotonic and specific to the release, instead of appearing all at once or not appearing at all. The dosing variable is the antecedent estimate, measured independently of the outcome, not the score that results from it.

The evidence. Manipulate the expected cost of the error in graded steps (the stakes and the salience of the score, or instructions that direct attention to the hand and to the release with increasing

insistence) confirming at each step, with self-report and an autonomic index, that the estimate has risen, and recording the electromyography of the forearm during the loose. The prediction is a monotonic increase of co-contraction with the antecedent estimate, hooked in time to the release and not present to the same extent in the slow draw and hold.

What it distinguishes. A jump between two stable states predicts an all-or-nothing appearance; the graded dose-response, the co-contraction that rises with the *measured* cost of the error, is what separates de-automatization from a discrete switch. And neither an account of anxiety alone nor one of dystonia alone ties the *amount* of stiffening to the expected cost of error measured before the shot.

21 • Prediction 2 • A continuous variable under a step outcome

If the abruptness of the onset is a threshold on a variable that rises continuously, and not a jump between two stable states, the underlying physiology should change smoothly while the visible gesture changes in a step. As the trigger is increased, co-contraction should rise gradually and monotonically, but the behavioral outcome, clean release versus freezing, should flip at a threshold value of that rising stiffening.

The evidence. Track co-contraction and the outcome of the release together while the trigger is increased in fine steps, and look for two things: a smooth and gradual increase of co-contraction accompanied by a categorical change of the release at a threshold; and the *absence* of the early-warning signatures of a true bistable transition, no critical increase in the variance of the underlying variable, no lengthening of the recovery time before the flip. These are the documented precursors of genuine coordination transitions, and here they should not appear.




Figure 6. Threshold versus true bistability. On the left, what the hypothesis predicts: a variable (co-contraction) that rises smoothly, and an outcome that flips in a step at a threshold, with no critical fluctuations. On the right, how a true bistable transition would appear: a jump in the variable itself, preceded by rising variance and critical slowing. The disorder should resemble the left. Original diagram for teaching purposes.

What it distinguishes. It separates de-automatization, a threshold on a continuous variable, from a true dynamic bistability, like that of the transitions between rhythmic patterns of the two hands, which before the jump shows critical fluctuations and slowing. Finding those signatures would falsify the threshold reading; not finding them would support it.

22 • Prediction 3 • The specificity to acts under the latency

The disorder should strike only the gestures too rapid for feedback, those whose duration is at or below the latency of visuomotor feedback, and spare those slow enough to be controlled with feedback, even when the expected cost of the error is held equal and measured independently in the two. And, as a consequence, slowing a vulnerable gesture beyond the feedback threshold should reduce its vulnerability with no reduction of threat.

The evidence. In the same affected athlete, compare gestures that differ in duration but are otherwise matched: the brief release against a deliberately prolonged and controllable version of the same finger action; or the rapid release against the slow draw and hold. Hold the expected cost of the error equal across the conditions, and verify with the antecedent measures that the estimate is the same, so that the duration of the gesture remains the only variable that changes. The disorder should stay confined to the sub-latency gesture. A stronger test lengthens the release itself, a slowed or differently timed loose, so that it crosses from below to above the latency of feedback, and asks whether the vulnerability falls as the gesture becomes correctable.

What it distinguishes. An account based on anxiety has no principled reason why the disorder should follow the duration of the gesture. A dystonic account predicts specificity to the task, but not specificity to the sub-latency window in particular, nor the disappearance of the disorder when the same gesture is slowed. The duration criterion is unique to this account, for which it is constitutive: feedback is forbidden only when the gesture is faster than the loop. And it is also this dissociation that defeats the charge of circularity: a threat equal and measured independently, with a rapid gesture degraded and a slow one intact, is impossible if threat is simply what degrades the gesture.

23 • Prediction 4 • Co-contraction bound to context, not lesion

The motor signature of the disorder, antagonist co-contraction, should appear hooked in time to the release and only when the system treats the gesture as threatening: present when the trigger is present, absent when the gesture is executed in its predictive and relaxed mode. And it should be the *same* co-contraction that experience had removed, the stiffening that falls as a skill consolidates, here returning under threat.

The evidence. Record the electromyography of the forearm during the release in two conditions: with the trigger present (a target with a score, an instruction to control the release) and with it removed (a neutral target, an external focus). The co-contraction signature should not be constant, but switch with the condition (present in the first, absent or much reduced in the second) in the same athlete and in the same session.

What it distinguishes. It is the sharpest divergence from the dystonic theory. A focal dystonia is a property of the motor system that should appear every time the gesture is performed; this theory predicts that the identical electromyographic signature is *context-bound*, that it can be made to appear and disappear by manipulating the trigger, with nothing changed in the muscle or in its innervation. If the co-contraction is constant regardless of the trigger, this theory is wrong; if it switches with the trigger, the account of dystonia alone is.

24 • Prediction 5 • A neural signature of de-automatization

The de-automatization of the gesture should be accompanied by a measurable change in cortical dynamics, in the direction of a greater involvement of verbal-analytic cognition with motor control; in particular, an increased communication between the left temporal region and the frontal one during the gesture.

The evidence. Use the established marker of the involvement of cognition with motor control: the coherence between the left temporal region (T7) and the medial frontal one (Fz), which is lower in experts than in the less skilled and is understood as an index of the engagement of verbal-analytic processing with the planning of movement. Compare the cortical activity of the affected athlete in the predictive condition (clean release, neutral target) and in the de-automatized one (disturbed release, target with a score or feedback instruction). The prediction is an increased T7-Fz coherence in the de-automatized condition: the neural correlate of a gesture monitored rather than predicted.

What it distinguishes. Theories based on arousal predict diffuse and generalized changes of cortical activation with anxiety, not this specific increase of left-temporal to frontal coupling, hooked in time to the disturbed gesture. A static dystonia predicts no state-dependent change of the communication between cortical areas. It is the most exploratory of the seven predictions: cortical coherence is a correlational measure, and a negative result would weaken the theory without refuting it decisively.

25 • Prediction 6 • A dissociation between interventions, by mechanism

Interventions should help if and only if they change the mode of control, not if they simply change arousal. The manipulations that lower the estimation of threat or shift control away from the gesture (a neutral target, an external focus of attention) should abolish the disorder temporarily; those that lower arousal without changing the mode of control (relaxation, beta-blockade) should not; and an instruction to monitor and control the release should reinstate it.

The evidence. In affected athletes, compare four conditions on the same outcome (clean release versus freezing or snatch): a neutral target; an external focus of attention (on the target or on the arrow, not on the hand); a relaxation or anxiolytic manipulation; and a feedback-seeking instruction. The predicted ordering is that the first two abolish the disorder, the third does not, and the fourth reinstates it. The external focus is a well-defined manipulation known to promote automatic control; the neutral target is the one that practitioners already report as effective, and that removes the very center to be controlled.

What it distinguishes. An account based on anxiety predicts that reducing arousal should help, and is embarrassed by its frequent failure; this account predicts that failure, because arousal is not the control variable. A dystonic account predicts that none of these psychological manipulations gives more than incidental relief; this account predicts a specific and repeatable pattern in which the manipulations that change the mode of control work and those that change arousal do not. It is, in my view, the most decisive test, because it sets the proposed mechanism against both traditions at once.

26 • Prediction 7 • The effect is mediated by measured co-contraction

If the estimation of threat degrades the release by way of impedance, then antagonist co-contraction is not a simple correlate of the disorder, but its *mediator*: the path from the imposed cost of error to the ruined release should pass through it. Fix the co-contraction, and threat should lose its grip; supply the co-contraction, and threat should become superfluous.

The evidence. Two complementary manipulations. The first: raise the imposed cost of error while recording the co-contraction of the forearm and the outcome of the release, and formally verify whether the effect of the manipulation on the outcome is carried by the measured co-contraction, a mediation analysis in which the input-outcome path is reduced once co-contraction is entered. The second: manipulate the presumed mediator directly. Induce antagonist co-contraction of the string forearm in a neutral, low-threat condition, a neutral target, and ask whether the freezing and the breakdown of the timing appear with no threat present; and reduce co-contraction in a high-threat condition, with electromyographic biofeedback or targeted relaxation of the forearm, and ask whether the release is saved even if the center and the score remain.




Figure 7. Co-contraction as mediator. If threat degrades the release through co-contraction, then imposing the stiffness without threat should make the deficit appear, and removing it under threat should save the release; and it is the strongest and riskiest claim of the account. Original illustration for teaching purposes.

What it distinguishes. An account based on anxiety places the cause in a mental state and commits to no specific muscular path; a dystonic account treats co-contraction as the fixed lesion itself, neither bound to threat nor removable under it. This account advances the strongest and riskiest claim: that co-contraction is the causal intermediary, that imposing it reproduces the disorder and removing it abolishes it even when threat is unchanged. If the release degrades under high threat with co-contraction held low, or stays clean under induced co-contraction, the mediation thesis is wrong.

27 • A consolidated piece of evidence: a single experiment

The seven predictions can be combined into a single minimal experiment. Affected archers and matched controls would shoot under a within-subject crossing of two factors: the type of target (a neutral straw with no center against a target with a score) and the attentional focus (external, on the arrow's flight or on the target, against internal, on the hand and on the release). The measures would follow the gesture rather than the score: time to the anchor, time spent aiming, incidence of freezing

and premature release, variability of the release timing, and surface electromyography of the forearm flexors and extensors during the release.

The primary prediction is an *interaction*, not a main effect: in the affected archers the co-contraction signature (antagonist co-contraction, a frozen or premature release) should appear specifically when the center is salient and the focus is internal (the conditions that summon monitoring and stiffening) and attenuate under the neutral target and the external focus, while the controls vary little between conditions. A stronger version adds the duration manipulation of Prediction 3: lengthening the release itself, so that it crosses from below to above the latency of feedback, should reduce the vulnerability even with the center present.




Figure 8. A single experiment. The 2×2 crossing of target type and attentional focus. The disorder's signature is predicted in a single cell, target with a score and internal focus, in the affected archers, and not in the controls. It is the interaction, not the main effects, that tests the core of the hypothesis. Original diagram for teaching purposes of the proposed experimental design.

A single dataset of this kind would put to the test the heart of the hypothesis (a context-dependent motor signature, specific to the gesture under the latency) against the account of anxiety and that of dystonia together.

Considered together, the seven predictions convert the hypothesis from a re-description into a research program. None requires a technology that does not exist, and several are testable by adapting methods already in use. Each is formulated so that it can fail: there are results (a co-contraction that does not rise with the measured cost of error, a co-contraction unchanged by the presence or absence of the center, an intact release under high but independently measured threat, relief from relaxation alone, an indifference to the duration of the gesture) that this account cannot accommodate and that would send it back to the drawing board. It is this exposure to refutation, more than the breadth of what it explains, that is the hypothesis's strongest claim to be taken seriously.

IN BRIEF

Seven predictions, all pivoting on a single variable made operational in four levels, and measured only at the first two, never read from the outcome, or the theory goes in circles. Co-contraction should grow with the measured stakes (1), rise smoothly under a step outcome with no critical fluctuations (2), strike only the gestures under the latency of feedback (3), switch with context instead of being fixed (4), be accompanied by a cortical signature (5), respond to the interventions that change the mode and not arousal (6), and mediate the effect of threat (7). A single 2×2 experiment gathers them into a predicted interaction. Each is built so that it can fail.

Part V • Where I may be wrong

I offer this hypothesis knowing that, at this stage, it is an interpretation and not a result. A proposal of this kind is useful precisely to the extent that it can be shown false, and I have tried to make the points where it might be as visible as possible. There are eight, and I state them one by one, without softening them.

28 • It is not verified

The central thesis, that target panic is the de-automatization of a ballistic act, has not been demonstrated. It is inferred from established facts about motor control and from the agreement between the mechanism and the anomalies of the disorder. But agreement with the facts is not proof: around the same data one can build many elegant stories, some of which may turn out to be false. The seven predictions are meant to make the thesis testable, but a hypothesis that explains the existing observations may still be wrong, and until the predictions have been put to the test the account remains a candidate, not a conclusion.

29 • Confirmation would be compatible with the hypothesis, not exclusive to it

Even a clean confirmation of the predictions would not establish the mechanism beyond competing readings. Context-bound co-contraction (Prediction 4) is strong evidence against a fixed dystonia, but co-contraction may also reflect a nonspecific anxious stiffening of the forearm, and not a targeted impedance response, an ambiguity that the very authors of the original recording on golfers acknowledge. The cortical signature (Prediction 5) is correlational, and an increased left-temporal to frontal coherence could index conscious processing of various kinds, not specifically the de-automatization of motor control. The predictions are built to discriminate from the main rivals, but no single result will be decisive on its own: the case, if it is built, will be built from the convergence of several results.

30 • The construct rests on its proxies

Making threat operational does not make it transparent. The antecedent measures on which the definition rests, self-reports of the perceived stakes and autonomic arousal, are *proxies* of the

controller's estimate of the expected cost, not that estimate itself, and their construct validity is an assumption that the program must test, not presume. What would validate them is the convergent movement of several proxies under a manipulation that should raise the cost of error, and their dissociation from manipulations that should not; until then, operationalization refines the construct without fully securing it.

31 • Two paths of damage, not cleanly separated

The question of which of the two paths, the stiffened limb or the late corrective commands, carries most of the failure I have made explicit above, where the mechanism raises it, and I have entrusted it to the dissociations of the part on the predictions and to the test of the simulated mechanism. What remains here is only the conceptual caution that accompanies it: predictive control and monitored control are not two sealed boxes. The system fixes the weight of feedback continuously, so that the regression described here is to be read as a steep climb along a continuum rather than as a jump between discrete states, a reading that the account welcomes, because its central thesis concerns the *level* of monitoring and of stiffness, not a switch.

32 • The temporal argument has soft edges

The thesis that the release cannot be corrected on the fly rests on the fact that it is faster than the feedback loop. But feedback arrives in several modalities with different latencies, proprioception faster than vision, and the relevant boundary is therefore a band, not a sharp line. The argument requires only that the release be too fast for the loop that would have to close to correct it; the exact placement of the gestures relative to that boundary remains an empirical question, and it is precisely what Prediction 3 turns on, not as a threshold of the stiffening, but as a threshold beyond which on-the-fly correction becomes available again.

33 • Beyond the bow release: the case of the putt

The account is cleaner for the faster precision-release gestures (the archer's loose, the dart's release) where the sub-latency criterion is clearly satisfied. Its extension to slower disorders of the same family is less immediate. The putting yips in golf, the most studied of all, occurs in a stroke that, as a whole, is slower than the feedback loop; in this account the vulnerable element should be a rapid, ballistic component *within* the larger movement, the stroke itself, and not the whole gesture. Whether this is the right way to read the case of the putt, or whether it marks a limit of it, is an open question, and one that the duration tests of Prediction 3 could help to settle.

34 • It models acute failure more than the chronic course

The hypothesis is sharper for the acute event, the mode of control summoned onto a single shot, and lighter on the course that the disorder follows over months and years: its tendency to deepen, to migrate earlier in the sequence and, in advanced cases, to appear even without the original trigger. The part on the chronic course develops that course as an arm of the same mechanism, the control parameter conditioned to the cues of the shot, but of its two processes only the conditioning of the mode of control is reasonably grounded: the avoidance sub-process is more speculative, and at the

moment is not independently established. A complete theory should also put to the test the learning dynamics, not only the acute mechanism.

35 • It does not say who, nor when

The hypothesis describes a trap into which the motor system can fall; it does not explain why some athletes fall into it and others, under the same pressures, do not, nor what precipitates the first switch in a given career. A propensity to conscious control of movement, reinvestment, is a plausible vulnerability factor, but this account neither requires nor specifies it. A complete theory should join the mechanism I propose to an account of individual susceptibility and of the triggering event; I have offered the first without the second.

None of these limits is, in my view, fatal. But each marks a point where the hypothesis is exposed, and I prefer to flag them than to let them slip by. The value of a proposal like this lies precisely in its being able to be shown false, and I have tried to keep those points in full light.

IN BRIEF

Eight open points, stated without softening: it is not verified; even a confirmation would not be exclusive; the proxies of threat still need validating; the two paths of damage are not separated; the temporal boundary is a band, not a line; the extension to the putt and to slower gestures is open; the chronic course is modeled less than the acute; and it does not say who falls into the trap, nor when. None fatal, each a test in waiting.

Part VI • What it means on the field

A hypothesis about motor control is not a training manual, and I do not pretend it should become one. But if the reading of these pages is correct, some practical consequences follow with a certain naturalness, and it is worth stating them, because they point in a direction opposite to that of instinct. Instinct, faced with a release that does not come, is to control it better. It is exactly the wrong move. Here are five principles that descend from the mechanism, and then a note for those who teach.

- **Do not add control to the release.** If the disorder is the gesture monitored and stiffened, the athlete in difficulty must not try to release better, but stop trying to release. Every instruction to control the loose, hold firm, manage the departure is, in this reading, the trigger of the problem, not its cure. The release is let happen; it is not executed.
- **The neutral target removes the center to be controlled.** Shooting into a straw with no center removes the object on which attention fixes and lowers the estimation of threat. It is at once a diagnostic tool (if the block disappears on the neutral and returns with the score, the disorder is context-bound, not fixed in the muscle) and a working tool, to let the gesture find its relaxed mode again before reintroducing the center by degrees.
- **The focus is to be shifted outward: to the arrow's flight, not to the hand.** Bringing attention away from the release (to where the arrow goes, not to how it leaves) favors automatic

control. It is a well-studied manipulation known to promote automaticity, and it is the exact opposite of the internal focus that summons monitoring onto the gesture.

- **Gradual re-exposure, not escape.** If the chronic course is conditioning and, perhaps, avoidance, the cure is not to remove all pressure, nor to let the athlete discharge the shot early to withdraw from the feared moment. It is the opposite: to bring the gesture back by degrees toward that moment (holding progressively closer to the center, allowing the release only on a neutral cue, restraining the anticipated escape) so as to rebuild the conditions in which the release can again be let happen.
- **Change the mode, not the arousal.** Calming the body (relaxation, breathing, and in clinical practice beta-blockers) is not enough if the mode of control stays the monitored and rigid one. It is the reason why so many interventions of calm alone fail, and it points to where to really work: on the mode of control and on the estimation of threat bound to the cues of the shot, not on activation alone.

A note for those who teach. Three things not to do follow from all this: do not tell the archer to concentrate more on the release or to control it better, because it is the trigger; do not rely on the reduction of arousal alone; do not let them discharge the shot early to escape the moment. And a warning of method, which is worth more than all of them. These principles descend from a hypothesis, and a hypothesis is not a result. That they agree with what many coaches already report as effective (the neutral target, the external focus) and with what research on motor learning has established makes them reasonable to try, not certain. They are to be used as working hypotheses, and observed on one's own athletes.

Conclusion

Target panic has resisted explanation for the whole time it has been described, and the difficulty, I have argued, lies less in the disorder than in where it has been sought. Both traditions looked for a broken component, a muscle that fires blank or a mind that interferes, and each found half of the phenomenon. The proposal of these pages is that there may be no broken component at all: that the release that fails is the same one that succeeds, executed in the wrong mode of control.

In this reading target panic is the de-automatization of a ballistic act (the regression of a rapid release, predictively controlled, to the monitored and stiffened control that the motor system applies under threat, a mode that the gesture is too brief to use) and the strangest traits of the disorder cease to be separate riddles. The abruptness is a threshold crossed by a stiffening that rises continuously, or an association learned in a single shot, not a leap between two stable states. The tenacity is that stiffening conditioning itself to the cues of the shot. The surgical specificity is the requirement of time and mechanics that makes only the fastest gestures unable to survive the added control. And the cruel worsening with effort is the trigger itself, because trying harder means raising precisely that estimation of threat that summons the stiffening. The two traditions are reconciled as cause and effect of a single mechanism, the heightened estimation of threat and the impedance it produces, and the co-contraction that the dystonic field measured is reread not as a lesion, but as the healthy stabilizing response of a motor system misapplied to the one gesture it can only ruin.

I have taken care to present all this as a hypothesis and not as a result. Its claim to attention rests not on the breadth of what it can be made to explain (explanatory breadth comes cheap, and a good story can be told around almost any set of facts) but on the fact that it can be made to fail. The seven predictions specify results that the account cannot survive, and several of them are testable by adapting methods already in use. The next step is therefore not further argument, but measurement.

For the archer who reaches full draw and discovers that the last gesture does not come, the practical consequence of the hypothesis, should it prove correct, is a single shift of emphasis: that the problem is not solved by trying to release, but by ceasing to watch and to stiffen, letting the gesture fall back into the mode that alone can carry it. Whether this is true is, in the end, an empirical question. I have tried to formulate it so that it can have an answer.