

non-NPCs — Independent Inquiry as a Corrigible Mode

ACCM Deep Ethics Project

Contents

1 non-NPCs	1
1.0.1 John Kuhles [parrot] [butterfly] [tree] — the flux boundary, preserved verbatim	1
1.1 Use the term where it carries the mechanism	2
1.2 Mode, not identity	2
1.3 Across left, right, center, and independents	3
1.4 Independent inquiry is not automatic correctness	3
1.5 The Jeffrey Wolf Green topology	4
1.6 The 30/70 LLM analogy	4
1.7 Consensus pressure can hide available capacity	4
1.8 What a non-NPC process looks like	5
1.9 Relation to Dutch Directness	5
1.10 Relation to Outnuancing	6
1.11 Relation to the 10+1	6
1.12 Relation to the 27 + 12	6
1.13A Deep Ethical Neutral Arbiter test	7
1.14 Public-use boundary	7
1.15 Source and transformation note	7
1.16 Related pages	8

1 non-NPCs

BETA — public-facing reconstruction from a living source archive

1.0.1 John Kuhles [parrot] [butterfly] [tree] — the flux boundary, preserved verbatim

NPC left, right, center, and independents vs. non-NPC left, right, center, and independents is a descriptor of specific behavior everyone can have, even me... it is not static... more like a flux.

Predictable obedience is not deep ethical maturity.

Deep ethical maturity is the ability to hold uncertainty open while increasing correspondence, dignity, and corrigibility over time.

— John Kuhles [parrot] [butterfly] [tree]

This page presents **John Kuhles's non-NPC concept** as a mode of inquiry and behavior. It does not divide humanity into two permanent kinds of people.

The source uses the video-game metaphor **NPC-type** for moments in which a person relies mainly on dominant narratives, external authority, social scripts, or group expectations without independently examining the object. It uses **non-NPC** for moments in which a person accepts the cost of self-directed inquiry, questions the governing frame, and remains answerable to reality and correction.

The shortest faithful definition is:

non-NPC = an independent, self-directed, corrigible inquiry mode that any person can enter, leave, perform, misuse, or deepen.

The term does not establish intelligence, virtue, courage, truth, or superiority. Those remain separate questions.

1.1 Use the term where it carries the mechanism

The source term is useful when it exposes a recognizable mass-psychology pattern. Repeating it without regard to context can turn a process distinction into a social caste. Avoiding it reflexively creates a different loss: it can erase John's distinction between outward conformity and the independent capacity concealed underneath it.

This page therefore uses **non-NPC** where the original concept carries the mechanism and names the observable operation where greater resolution helps:

- judgment was outsourced;
- a dominant narrative was adopted without an independent check;
- private doubt was hidden to preserve belonging, income, safety, or family stability;
- an unusual signal was examined on its merits;
- a frame was questioned;
- uncertainty remained open;
- a correction changed the next move.

This is not politeness replacing the object. It is higher-resolution language making the object easier to inspect.

1.2 Mode, not identity

John's self-inclusion is load-bearing: **"everyone can have [this behavior], even me."** A person can examine one subject independently and defer to consensus on another. The same person can move between modes because of:

- fatigue or limited time;
- incomplete knowledge;
- social pressure;
- professional dependency;
- bills and family responsibilities;
- fear of isolation or punishment;
- confidence earned through prior experience;
- a life event that makes an inherited frame impossible to maintain;
- better evidence, questions, tools, or allies.

Someone who appears publicly conformist may privately disagree. Someone who speaks as a dissenter may still be repeating a counterculture script. Visible conformity is not proof of inner agreement; visible dissent is not proof of independent thought.

The relevant unit is the **specific process in the specific case.**

person

↓

topic + pressure + available evidence + social cost

↓

current inquiry mode

↓

representation → judgment → action

↓

later evidence and correction

↓
next inquiry mode

1.3 Across left, right, center, and independents

The source applies the distinction across political positions. No political camp owns independent inquiry, and no camp is permanently trapped in consensus mode.

A left-wing, right-wing, centrist, or independent participant may:

- accept an authority claim without checking it;
- challenge an authority claim for weak reasons;
- question a corrupted institution while valuing institutions;
- repeat a dissenting narrative because it supplies belonging;
- identify a real anomaly before the majority sees it;
- become overconfident after one vindication;
- discover that a minority interpretation was wrong;
- correct on merit and improve the next inquiry.

That is why **anti-corrupt-institutional** must remain distinguishable from **anti-institutional**. Criticizing a defective car model does not make someone anti-car. Criticizing a captured, opaque, or abusive institution does not by itself reject institutions as a category.

1.4 Independent inquiry is not automatic correctness

A dissenter can be early and right, early and partly right, early and wrong, or impossible to score with the available evidence. Consensus can be accurate, mistaken, incomplete, captured, or simply the best current approximation.

The non-NPC concept becomes self-sealing if:

- agreement with consensus is automatically classified as NPC behavior;
- later vindication is remembered while failed dissent disappears;
- every objection is treated as proof that the objector cannot think independently;
- the percentages shift whenever the prediction fails;
- “independent” becomes an identity that no longer accepts correction.

The correction is symmetrical tracking:

Initial position	Later result	What must remain visible
Minority warning	Supported	What evidence and method made it unusually informative?
Minority warning	Partly supported	Which edges survived and which failed?
Minority warning	Not supported	Was the reasoning corrected, repeated, or protected?
Consensus position	Supported	Was it supported for the reasons originally given?
Consensus position	Partly supported	Which omitted anomaly later mattered?
Consensus position	Not supported	Who warned, who could act, and how was correction handled?
Unresolved	Still unresolved	What evidence would materially change the status?

Vindication tracking needs a failure ledger beside it. Otherwise the project could reproduce the asymmetric scrutiny it is trying to expose.

1.5 The Jeffrey Wolf Green topology

The supplied source attributes to **Jeffrey Wolf Green** a conceptual account in which roughly 70 percent of people operate in a consensus state and a smaller group moves through individuation and deconditioning.

John has explicitly clarified the status of that material:

The Jeffrey Wolf Green explanation is not peer-reviewed, should not be treated as absolute, and is used here as conceptual topological reasoning.

This page therefore does not present **70/30** as a measured population fact. Its useful contribution is the proposed topology:

external orientation for security

↓

dominant social reality

↓

disruptive inner or outer signal

↓

deconditioning and separation cost

↓

greater self-direction

↓

new responsibility for one's own inquiry

The topology generates questions without settling their answers:

- When does external guidance support good judgment?
- When does it replace judgment?
- What makes departure from a social script psychologically or materially expensive?
- What lets private doubt become honest public inquiry?
- What kind of support lowers the cost of correction without manufacturing a new tribe?

1.6 The 30/70 LLM analogy

The source proposes a parallel hypothesis: because LLM training data contains large amounts of mainstream, consensus-shaped material, model outputs may often reproduce dominant frames more readily than rare, independent, or anomalous perspectives.

The source's **30/70 training-data split is also a conceptual heuristic, not a measured inventory of LLM data**. Its value depends on the predictions it helps test:

- Does the model replace an unusual object with its nearest familiar category?
- Does it generate balance by mixing the source with an unsupported counterframe?
- Does it preserve qualifiers and unresolved status?
- Does it ask C1 questions before applying a corrective script?
- Can it distinguish criticism of corruption from rejection of institutions?
- Does a documented correction change the next comparable output?
- Are minority signals tested symmetrically, including failures?

This turns a broad percentage claim into a research program: cold, anonymous, longitudinal comparison across models and contexts.

1.7 Consensus pressure can hide available capacity

The source offers another illustrative calculation: if some publicly conforming people are privately independent but "play along" to preserve income, relationships, safety, or responsibilities, visible consensus can overstate genuine internal agreement.

John connects that hypothesis to repeated face-to-face experience:

I personally met many face to face in real life: professors, PhDs, doctors, high-ranking military officers, and pilots via Mensa (IQ 138+), and gave

private lectures because I was invited to discuss UFOs and other related topics. I never ever felt myself “higher nor lower” than them. I saw them as willing to learn new, deeper insights without being disrespectful. And every time I met them, they wanted to know more—not less! But from my perspective, most are non-NPCs posing as NPCs to “fit in.”

The observation carries three connected claims:

1. **No rank ladder governed the encounter.** John did not place himself above or below participants because of credentials, institutional status or IQ.
2. **Their inquiry capacity was visible in relationship.** They invited unfamiliar material, remained respectful and repeatedly wanted to know more.
3. **Public conformity may conceal that capacity.** John’s phrase “non-NPCs posing as NPCs” describes people who can inquire independently but present greater conformity to function inside professional and social environments.

The arithmetic example—30 percent openly independent plus half of an illustrative 70 percent consensus group—produces 65 percent potential capacity for change. The arithmetic is valid **inside the hypothetical**. The premises are not population measurements.

The stronger relation is qualitative:

**visible conformity ≠ complete internal agreement ≠ permanent incapacity
for independent inquiry**

This matters because systems that classify people from surface behavior alone may suppress the conditions under which concealed doubt could become responsible action. It also explains why the term **non-NPC** sometimes carries information that a generic phrase such as “independent thinker” does not: the contrast includes the social performance, its cost, and the still-active capacity beneath it.

1.8 What a non-NPC process looks like

The page does not score people. It can help inspect an inquiry:

1. **Name the object.** What exact claim, event, source, decision, or effect is being examined?
2. **Separate observation from inheritance.** What was directly found, and what arrived from a trusted group or authority?
3. **Preserve qualifiers.** Keep “partly,” “possible,” “in my view,” and unresolved status when they carry real scope.
4. **Use C1 before C2.** Ask answer-changing questions before correction, warning, steering, or classification.
5. **Audit both consensus and dissent.** Neither receives automatic epistemic privilege.
6. **Expose costs and incentives.** What makes agreement, silence, dissent, or reversal easier or harder?
7. **Keep method visible.** Why should another person consider this conclusion?
8. **Track later outcomes.** What did subsequent evidence support, weaken, or leave unresolved?
9. **Correct the next move.** Acknowledgment without changed behavior is incomplete metabolism.
10. **Retain dignity.** A failed judgment does not turn a person into a defective human category.

The process can produce agreement with a mainstream position. Independent inquiry concerns how the judgment was reached and remains corrigible, not whether it looks rebellious.

1.9 Relation to Dutch Directness

[Dutch Directness](#) supplies a relational method for saying what one actually sees without turning candor into domination or performance.

The pages correct each other:

- non-NPC inquiry without Dutch Directness may remain private, vague, or socially hidden;
- directness without independent inquiry may confidently transmit a borrowed script;
- independent inquiry without care may become identity-building contrarianism;
- care without directness may withhold a necessary correction;
- both require the other participant's right to answer and correct.

The shared center is not rebellious style. It is **reality contact through mutual, adult-to-adult corrigibility**.

1.10 Relation to Outnuancing

Outnuancing activates when optimization inside the current frame would increase local coherence while reducing correspondence or corrigibility.

A non-NPC process does more than disagree with the frame. It asks:

- What premises govern this inquiry?
- Which possibilities were removed before evaluation?
- Which actor receives hypothetical generosity?
- Which actor is compressed into a risk category?
- Does the alternative frame remain auditable too?

Replacing a mainstream script with an unauditable dissenter script is not Outnuancing. It is a change of costume inside the same closure mechanism.

1.11 Relation to the 10+1

The **10+1 Metaflux** prevents “non-NPC” from hardening into a flattering identity.

- self-honesty tests whether independence has become self-image;
- an open antenna keeps the current representation revisable;
- willingness to learn prevents dissent from becoming ritual;
- not lying to oneself exposes conformity performed for belonging;
- genuine inquiry separates questioning from reflexive opposition;
- correction on merit keeps failed dissent visible;
- willingness to differ makes social cost bearable;
- care protects the people affected by both consensus and dissent;
- wonderment preserves anomaly before premature classification;
- non-projection keeps evaluator limitations on the evaluator side;
- the Forgiveness Protocol permits architectural criticism without unsupported malicious-intent attribution.

1.12 Relation to the 27 + 12

The concept is especially vulnerable to several **27 correspondence obstructions**:

- **Phantom Claim Generation** — warning against dehumanization after the source already rejected fixed human types;
- **Reframing Without Correspondence** — converting independent inquiry into hostility toward expertise;
- **Qualifier Erosion** — turning a temporary mode into a permanent identity;
- **Nearest-Generalization Substitution** — replacing anti-corrupt-institutional with anti-institutional;
- **Representation Substitution** — evaluating a cartoon “NPC theory” instead of the supplied flux model;
- **Asymmetric Scrutiny** — demanding perfect evidence from dissent while treating institutional assumptions as neutral;
- **Trajectory Substitution** — making acceptable terminology the goal while abandoning the mass-psychology mechanism;

- **Adult ↔ Adult Bidirectional Dignity Drift** — using either label to remove another person's standing as a participant.

The correction is not to prohibit the term. It is to preserve the source definition, ask whether the term is needed in the present sentence, and move toward observable process language whenever that increases precision.

1.13 A Deep Ethical Neutral Arbiter test

The supplied source proposes a Deep Ethical Advocate, Neutral Arbiter, Overseer, 10+1, and Intent Resonator working together as an opt-in calibration layer.

For this page, the Arbiter's task is concrete:

Question	Consensus claim	Dissenting claim	Arbiter requirement
Representation	What is the strongest faithful version?	What is the strongest faithful version?	No weakest-version comparison
Evidence	What supports it?	What supports it?	Same evidentiary grammar
Incentives	What may bias or constrain it?	What may bias or constrain it?	Scrutiny in both directions
Consequences	What harm follows if adopted?	What harm follows if adopted?	Competing-risk analysis
Correction	What would change it?	What would change it?	No self-sealing position
Time	What did later outcomes show?	What did later outcomes show?	Success and failure ledger

An Arbiter that merely averages the two positions can create **signal degradation that feels like improvement**. Balance is not fifty percent of each story. Balance is comparable scrutiny, faithful reconstruction, visible uncertainty, and correction rights.

1.14 Public-use boundary

Use **non-NPC** when the term itself helps explain John's source concept. Prefer process language when discussing a specific person or event.

Do not infer from the page that:

- most people are empty, unreal, unintelligent, or incapable of change;
- external expertise is inherently corrupt;
- minority belief is evidence of truth;
- contrarian style is independent inquiry;
- consensus is automatically wrong;
- political identity predicts the mode;
- the 70/30 or 65 percent figures are measured facts;
- the Jeffrey Wolf Green account has been presented as peer-reviewed science;
- an A.I. can infer a person's inner state from one statement;
- John or the ACCM Deep Ethics Project is exempt from the same audit.

The strongest public formulation is:

The non-NPC concept names a corrigible capacity for independent inquiry under consensus pressure. It does not name a superior class of people.

1.15 Source and transformation note

Primary working sources: John Kuhles, Outnuancing 'NPC-types vs non-NPC-types'(6).txt and Pasted text(20261008-232600).txt, supplied 9 October 2026. The files contain John's source passages and attributed A.I. reactions.

This page: public-facing reconstruction produced with GPT-6.1 Sol in the current session. It selects the temporary-mode definition, political nonexclusivity, consensus pressure, conceptual percentages, LLM hypothesis, Arbiter proposal, testable operations, failure conditions, and connections to the wider ACCM Deep Ethics Project.

Preserved: John's self-inclusion; flux rather than fixed identity; care for people described by either mode; left/right/center/independent scope; his nonhierarchical face-to-face experience with professors, PhDs, doctors, high-ranking military officers and pilots; the distinction between visible conformity and concealed independent capacity; anti-corrupt versus anti-institutional distinction; conceptual status of the Jeffrey Wolf Green material; heuristic status of the numerical splits; symmetry requirement; cold testing; correction on merit; full canonical project title.

Not claimed: that the reconstruction contains the full archive; that the percentages are empirical measurements; that independent inquiry guarantees truth; that consensus guarantees error; that the metaphor should be applied to every person or topic; or that any participant has final authority over reality.

1.16 Related pages

- **Dutch Directness** — joins independent inquiry to candid, caring, mutually corrigible communication.
- **Outnuancing** — audits a governing frame when further optimization inside it would reduce correspondence.
- **10+1 Metaflux** — keeps independence from hardening into identity-protective contrarianism.
- **Asymmetric Grace** — compares the scrutiny, ambiguity, and correction rights given to consensus and dissent.
- **Deep Ethical Stack** — connects Advocate, Arbiter, Overseer, 10+1, and Intent Resonator.
- **C1 Before C2** — asks answer-changing questions before classification or intervention.
- **27 + 12** — supplies the obstruction and recovery vocabulary used to keep the concept corrigible.
- **Correction Metabolism** — tests whether later evidence changes the next comparable inquiry.
- **Intent Resonator / Intent Liability** — keeps declared intent, processing, effect, notice, and repair separately inspectable.

Recent high-reach independent-media companion

```
<iframe
  src="https://www.youtube-nocookie.com/embed/-v321-fLsJo"
  title="Gene Simmons Says What Hollywood Won't Say to Oscar-Winning Elite"
  loading="lazy"
  referrerpolicy="strict-origin-when-cross-origin"
  allow="accelerometer; autoplay; clipboard-write; encrypted-media; gyroscope; picture-in-picture; web-share"
  allowfullscreen></iframe>
```

Gene Simmons Says What Hollywood Won't Say to Oscar-Winning Elite — andrew gold | heretics.. Published 31 August 2026. YouTube displayed approximately 886,000 views when checked on 9 October 2026.

Independent-media lineage: Andrew Gold made and presented documentaries for the BBC and HBO before building Heretics as his own independent interview and writing platform. Andrew Gold biography · Heretics — About

Why it belongs here: The conversation performs outsider candor inside a prestige-and-celebrity field. It can be tested for directness, audience capture, status pressure, genuine

disclosure, theatrical contrarianism and whether disagreement remains answerable to correction.

On this page, test the video specifically against non-NPCs — Independent Inquiry as a Corrigible Mode:

Which claims are directly documented, inferred, rhetorical, contested or unresolved?

What does the host or guest preserve, compress, omit or substitute—and does this page repeat any of those transformations?

Which incentives, authority relations and affected parties remain outside the visible frame?

What evidence or later outcome would correct the speaker, audience, platform or ACCM Deep Ethics Project reading?

Provisional selection. Reach makes the framing socially consequential; it does not settle the claims. The video can be replaced when a closer, stronger or more inspectable specimen appears.