

THE IN A HURRY SERIES

Persuasion

in a Hurry

How minds are actually changed

Max Brocklesby

BOOKSINAHURRY.COM

Contents

The Whole Thing in One Page 4

Why You Should Care 5

The Core Ideas 7

How It Actually Works 20

What People Get Wrong 31

Use It 36

Terms 41

Go Deeper 44

Notes and Sources 47

Bibliography 54

The Whole Thing in One Page

Persuasion has been sold as a box of verbal lock picks. Find the right phrase, mirror the right gesture, mention scarcity, press the hidden button and another mind opens. That picture survives because it flatters the persuader. It makes failure a technical error and the other person a mechanism waiting to be operated.

A mind is not an empty container. It arrives with beliefs about what is true, attitudes about what is good, loyalties about who can be trusted, habits about what happens next and a sharp interest in remaining its own.

Persuasion begins inside that existing structure. The same argument can look decisive, irrelevant or insulting depending on who hears it, from whom, in what setting and at what cost to their identity.

That makes persuasion a process rather than a trick. A message competes for attention and has to make sense. Its source shapes how seriously the claim is taken; the claim can also change the listener's opinion of its source. Thoughts and feelings meet prior knowledge and the question, often unspoken, of what agreement would mean about oneself and one's group. These influences interact rather than wait their turn. Practical barriers may prevent an intention becoming action. Later experience can strengthen a change or undo it.

Each stage can move without the next. A person may remember an argument and reject it, believe a fact without changing an attitude, form an intention and never act, or agree aloud while remaining privately unmoved. Immediate assent is therefore weak evidence of a changed mind. The stronger tests come later, when the speaker has gone and a choice must be made.

Reasons matter, but they work through the thoughts they produce. Emotion matters because it marks danger, value and urgency, not because feeling defeats reason. Credibility matters, but expertise without trust can harden suspicion. Repetition can make a claim familiar, and familiarity can be mistaken for truth, but repetition also irritates and cannot erase every prior

commitment. Social proof can guide uncertain people, yet advertising an undesirable norm can accidentally teach that the behaviour is common. Every familiar principle is conditional. Its effect depends on the outcome, audience, timing, message, source and comparison condition against which it is tested.

Agency complicates the whole process. Pressure, certainty and commands can make a threatened person search for objections and restore freedom by doing the opposite. Respectful questions, self-generated reasons, credible evidence and direct experience can make a new position feel owned rather than imposed. That is why listening can improve the conditions for persuasion, why a conversation can outperform a polished speech, and why a person may change after saying the argument in their own words.

The ethical line follows the mechanism. Persuasion offers reasons while leaving refusal possible. Manipulation hides relevant aims or exploits vulnerabilities. Coercion changes the consequences of saying no. All can change behaviour; they do not produce the same kind of consent.

Minds change when evidence, value, trust, social permission, feasibility and agency line up well enough for revision to become easier than resistance. The persuader can arrange those conditions. The recipient still does the changing.

That is the book.

Why You Should Care

You can win an argument and fail to persuade. The other person stops objecting, says they see your point and leaves. Ten minutes later they explain to somebody else why you were wrong. What changed in the room was the conversation, not the mind.

That distinction matters wherever another person has a choice. A doctor can explain a treatment accurately and still fail to make it acceptable. A manager can announce a sound change and create a rebellion by making it feel imposed. A campaign can move a poll answer for an afternoon and leave

the vote untouched. A parent can secure obedience while teaching a child to hide disagreement. In each case the words may be clear. The failure lies somewhere else in the chain.

Consider a hypothetical manager celebrating a successful presentation. Everyone could explain the new procedure. Almost everyone said it made sense. A month later, the old spreadsheet is still doing the work. The meeting may have taught something without changing what anyone does. Persuasion research separates those outcomes instead of letting a favourable response stand in for the result the manager wanted. An easy measure can flatter the speaker while the intended change never happens.

Once you see the sequence, everyday mysteries become legible. More evidence can fail when distrust is the bottleneck. A weaker argument can win because it lets the listener preserve dignity. A public commitment can alter later conduct, but it can also produce a performance for the room. Telling people that a bad behaviour is widespread can normalise it. A frightening warning can motivate, especially when it includes a credible response; dread alone tells nobody what to do. The outcome depends less on the existence of a technique than on what the technique does inside a particular audience.

This also protects you from the persuasion industry. The market is full of universal claims because universality sells better than moderation. Seven principles, three words, one body-language ratio, a colour that closes deals, a question that bypasses resistance. Real findings are often hiding underneath, but they have been stretched from a measured average in a defined setting into a promise about everybody. The useful lesson is rarely that the effect is fake. It is that the conditions are the lesson.

The subject matters for defence as much as influence. Familiarity can feel like evidence. Group agreement can stand in for private inspection. A confident source can borrow authority from clothing, status or fluency. Digital systems can test thousands of messages, select the ones that move measured responses and personalise them at low additional cost. Conversational artificial intelligence adds a persuader that can listen, adapt

and continue without tiring. None of that creates mind control. It does make poor mental models expensive.

There are limits. People are not clay, and even well-designed attempts often produce small, temporary or uneven effects. Many celebrated studies use students, brief exposures and self-reported attitudes. Politics, health, consumption and close relationships do not share one average effect. A message that works in one culture or identity group may misfire in another. This book will not give you a script that defeats those facts.

It will give you a better question. Instead of asking which trick changes minds, ask what must become true for this person to reconsider, and which part of the path from hearing to lasting action is blocked. That turns persuasion from performance into diagnosis. It also makes ethical failure harder to disguise, because the same model reveals when influence depends on hidden intent, manufactured pressure or the removal of a real choice.

The aim is not to become irresistible. It is to understand why resistance exists, when revision is possible, and how to make a case another person can accept without surrendering ownership of their mind.

The Core Ideas

First Decide What “Changed” Means

Before choosing a message, decide what would count as success. Persuasion can alter what a person notices, understands, believes, values, intends, says, chooses or keeps doing. These outcomes are related, but none is a reliable substitute for the next.

Imagine being offered a free trial of a service you do not much like. You accept because it costs nothing, use it once, then forget to cancel. Its owner can record a sign-up, a use and a paid renewal. None tells them whether your opinion improved. Equally, you might admire the service and decline because you cannot afford it. The phrase *changed their mind* hides these

differences in one flattering blur.

The Second World War supplied an early demonstration. The United States Army used a wartime film series to explain the conflict and strengthen soldiers' commitment. Carl Hovland and colleagues tested parts of the programme. The films could teach factual material more reliably than they changed broad attitudes or motivation. Information had entered. The desired conduct had not necessarily followed. The research did not show that facts were useless. It showed that learning a fact and reorganising a life around it are different events.

A useful outcome ladder runs from exposure through attention, comprehension, belief, attitude, intention, action and persistence. The order is not a law of the mind. People sometimes act first and revise their attitudes later. Habit can carry behaviour while conscious approval weakens. The ladder is a discipline for measurement: name the rung you need, then test that rung rather than the easiest one to observe.

A study showing immediate agreement does not by itself establish private conviction, behaviour next week or resistance to a counterargument. Recall may prove that a slogan stuck while saying nothing about whether it was believed. A statistically clear effect can still be the wrong effect.

Causation matters too. Agreement after a speech does not show that the speech produced it. The audience may have been sympathetic before arriving. An event outside the room may have moved everyone. People may tell the speaker what politeness requires. Strong evaluation asks what happened relative to a credible version of the world in which the persuasive attempt was absent or different.

Duration completes the definition. Some changes are meant to last only until a danger passes. Others must survive delay, social pressure, inconvenience and new evidence. The longer the claim, the stronger the test should be. Measure when the speaker has gone, refusal is safe and the relevant choice has become real.

The first question is therefore not *What should I say?* It is *What, exactly,*

should be different afterwards, for whom, compared with what, and for how long? Until that is answered, persuasion has no target and success has no meaning.

The Audience Is the Starting Condition

A persuasive message does not begin when somebody speaks. It begins with what the listener already carries into the room.

That starting condition includes beliefs, but beliefs are only one layer. People also have values, goals, identities, habits, memories of similar claims and expectations about the speaker. Some attitudes are casual preferences; others organise friendships, work, religion or political belonging. Some come from hearsay; others were learned through direct experience. The stronger, more accessible and more rehearsed an attitude is, the more likely it is to guide later judgement and behaviour. Calling all of this bias misses the structure. A person who ignored everything already learned would be impossible to educate and easy to exploit.

New information is interpreted relative to that structure. Suppose two people hear that a proposed rule will reduce risk. One values security and hears protection. The other values independence and hears control. The fact has not changed, but its meaning has. A message can fail because the evidence is weak, because the evidence answers a question the audience is not asking, or because accepting it appears to threaten something more important than the claim itself.

Identity raises the stakes. A position can become a badge showing what kind of person someone is and which people are theirs. Revision then carries two costs. The person must concede a factual or moral error, and may also risk disloyalty. More data does not remove that second cost. Sometimes it increases it by making the challenge harder to ignore.

This does not mean people reason only as tribal lawyers. They care about accuracy, social acceptance and a coherent sense of self in changing proportions. The same person can inspect a pension comparison carefully, follow friends into a restaurant and defend a political belief with remarkable

creativity. Context decides which motive becomes prominent.

The practical consequence is audience diagnosis. Before choosing an argument, ask what judgement is being requested, what the present position does for the person and what accepting the new position would cost. The answer may concern money or risk. It may concern dignity, status, loyalty or a previous public commitment. A technically correct message aimed at the wrong motive is still a poor message.

Matching can help. In a large 2022 synthesis, messages aligned with recipients' active goals, needs or values produced a positive average persuasive effect across hundreds of experimental tests. The result was not a licence for crude personality targeting. Effects varied, and an obvious or clumsy match can feel manipulative. It established a more modest claim: relevance is part of evidence as experienced by a human recipient.

Political studies of moral reframing make the mechanism visible. Advocates tend to offer the moral reasons that persuaded them, then seem surprised when opponents remain cold. In United States experiments, arguments framed around the recipient's moral values sometimes moved opinion more effectively. There is no need to counterfeit a conviction. Suppose a town is considering flood defences: protecting vulnerable residents and preserving inherited homes may both be honest reasons for the same project. Different listeners can arrive through different doors. The ethical requirement is that each door leads to the same proposal, with the same costs disclosed.

Persuasion therefore starts with an asymmetry. You know what you want to say. The listener knows what would make it matter. Until those meet, eloquence is decoration.

Attention and Understanding Open the Gate

For a reason to persuade, someone has to encounter it and grasp what it says. That is a lower bar than agreement. Bare exposure or a visual cue can influence judgement without a clearly understood argument; neither shows that the case itself has been evaluated. Communicators often mistake securing attention for finishing the job.

Attention is selective. People turn towards change, danger, novelty, reward, social information and whatever their current goals make relevant. They also miss material placed directly in front of them when another task consumes the same limited capacity. The practical contest is therefore not between your message and silence. It is between your message and every other claim on the recipient's mind.

Salience can be earned by consequence or manufactured by presentation. A red warning, a startling image or a personalised subject line may win the first second. If the claim beneath it does not repay attention, the device teaches avoidance. Attention tactics also select an audience: the feature that attracts one group can tell another that the message was not made for them.

Clarity earns processing time. Concrete language, a visible structure and one claim at a time reduce the work required to build meaning. Unnecessary complexity gives the listener two chances to reject you: first because the message is hard to follow, then because confusion can be attributed to weakness in the case. Yet ease is not truth. A polished falsehood can feel more plausible because fluent material is easier to process. A difficult truth may demand effort because the subject is difficult.

Repetition sits inside that tension. Repeated exposure can increase familiarity, and familiarity can be misread as evidence that a statement has been encountered because it is true. Experiments on the illusory truth effect show that prior exposure can raise later truth ratings even when people possess relevant knowledge. The effect is neither irresistible nor unlimited. Strong contrary knowledge, distrust, identity and context can blunt it, while excessive repetition can create boredom or suspicion. Repetition changes one input into judgement; it does not replace judgement.

The same distinction applies to memory. People can remember slogans they detest. A campaign may improve recognition of a candidate while also reminding voters why they dislike that candidate. Familiarity is often helpful when the main barrier is uncertainty or obscurity. It is dangerous to treat recall as endorsement.

Comprehension has layers. A listener may understand the words but miss the causal model. Saying that a treatment cuts relative risk by half may be grammatically clear while hiding whether the baseline risk fell from two in a thousand to one, or from two in ten to one. A graph can look transparent while changing its scale. Good persuasion makes the comparison inspectable. Bad persuasion exploits the difference between a sentence being easy to repeat and a claim being easy to evaluate.

There is also a timing problem. Complex arguments require enough motivation and capacity to think. Fatigue, distraction, time pressure and emotional arousal can narrow what gets processed. Under those conditions people may rely more on visible cues such as source, consensus or familiarity. That does not make them irrational. It makes them economical. No one can re-investigate every claim from first principles.

The first operational job is therefore to secure a fair hearing: reach attention, make the proposition clear, show why it is relevant and give the recipient enough cognitive room to inspect it. None of those establishes agreement. They open the court.

Credibility Determines Which Claims Get a Hearing

A claim arrives attached to a claimant. The listener judges both.

Modern persuasion research often divides credibility into expertise and trustworthiness. Expertise asks whether the source is likely to know. Trustworthiness asks whether the source is likely to tell the truth as they see it. A third quality matters in close interaction: goodwill, the belief that the speaker understands and cares about what happens to the listener. A brilliant expert who seems contemptuous may lose to a less qualified person who has earned relational trust.

None of these qualities belongs permanently to the speaker. Credibility is granted for a domain, by an audience, in a setting. A Nobel Prize does not confer expertise in pension design. A close friend may be trusted about motives and distrusted about medicine. An institution can possess excellent

data and a history that makes part of the public unwilling to accept its interpretation. Telling that audience to trust the institution because it is authoritative restates the problem.

The Yale research programme led by Carl Hovland after the Second World War made source, message and audience separate experimental questions. Later work has repeatedly found that credible sources can be more persuasive on average, while also showing qualifications. Effects depend on the issue, the timing of measurement, prior attitudes, the strength of the argument and whether the source cue is considered before or after the message. A source can shape how much thought occurs, supply evidence, bias the direction of thought or alter confidence in thoughts already formed.

That last mechanism explains a common surprise. The same strong argument may land differently when attributed to a respected specialist because the recipient becomes more confident in favourable reactions. The source did not bypass thought. It changed how the person evaluated their own thought. Conversely, an untrusted source can contaminate a claim so thoroughly that even accurate information becomes harder to accept.

Shared identity can bridge this gap. People often grant an in-group messenger more knowledge of their circumstances and less hostile intent. Yet identity is not a universal key. An insider who appears to betray the group may be judged more harshly than an outsider. Visible targeting can also expose the attempt: this message was selected to operate on people like me. Relevance becomes surveillance.

Credibility can be built through conduct. State what is known, what remains uncertain and what interest you have in the outcome. A mechanic who shows you the worn part offers something a confident voice cannot: a claim you can inspect. Make predictions specific enough to check later. Admit a limitation before an opponent discovers it. None guarantees trust, but it gives the listener grounds for judgement. Empty certainty spends credibility for emphasis and leaves none for correction.

Presentation still matters, but its role is often misunderstood. Fluency, confident delivery, professional design and social status can act as rough cues to competence. They can also be staged. The useful question is not whether appearance matters. It is whether the cue is connected to the knowledge and honesty required for this claim.

When trust is the bottleneck, adding another chart may achieve nothing. Diagnose the source problem directly: do they doubt that you know, that you are candid, or that you understand their interests? Those are different failures. Each needs different evidence.

The Listener Writes Half the Argument

You supply the words. The listener supplies much of what gives them force: remembered examples, suspected motives, a comparison you never mentioned. The argument taking shape in their head may be better than yours. It may also be an argument against you.

A sentence can produce support, objections or no sustained thought at all. Counting the claims you delivered misses what matters: the responses the listener formed, and how much confidence they place in them.

This is the durable insight behind dual-process models such as the elaboration likelihood model. When people are motivated and able to examine a relevant issue, the quality of the case matters more. They compare evidence, test implications and connect the proposal to what they know. Attitudes formed or changed through substantial thought tend, under many conditions, to be more stable, more resistant and more predictive of behaviour. When motivation or capacity is low, people may rely more on manageable cues: the source seems credible, the position looks popular, the message feels familiar, the argument count is large.

The popular version turns this into two roads, one rational and one gullible. That is too neat. A cue can become an argument. If experts have independently reached a strong consensus on a technical question, their agreement is relevant evidence for a non-specialist. Emotion can motivate deeper inspection. A weak source can make a listener scrutinise the case

and discover that it is sound. The same feature can affect how much thought occurs, which direction it takes and how much confidence the person places in the result.

Argument quality is also audience-dependent without being arbitrary. Evidence is strong when it bears on the proposition, survives reasonable alternatives and matters to the recipient's decision. A well-designed study on short-term preference may be weak evidence for lasting behaviour. Ten reasons are worse than two if eight invite easy rebuttal. A concession can strengthen a message when it proves the speaker has inspected the other side, though a concession that destroys the conclusion is merely honesty arriving on time.

Framing changes which valid aspect becomes prominent. A medical outcome can be described in survival or mortality terms. A policy can be presented as preventing a loss or securing a gain. These descriptions may be logically related while feeling different because people attend to reference points, responsibility and emotional meaning. Framing becomes deception when it withholds the denominator, comparison or consequence needed for judgement. A frame is unavoidable; a hidden frame is optional.

Stories work by supplying organised experience. They give a claim agents, causes and consequences that can be mentally simulated. Meta-analytic evidence finds modest average movement in narrative-consistent beliefs, attitudes and intentions, with much variation. That does not show that stories beat statistics. A vivid case can reveal a mechanism or distort a base rate. The strongest use joins human consequence to representative evidence, so the story tells the reader what the number means without pretending to tell them how common it is.

The recipient's own reasons can carry unusual weight. Questions that prompt someone to retrieve examples, compare alternatives or state why a change matters can produce self-persuasion. The person now supplies arguments selected from their own knowledge and expressed in their own language. A well-chosen question can therefore do more than another declaration, provided it invites thought rather than steers towards a

compulsory answer.

A persuasive case should therefore be designed for the thoughts it invites. Anticipate the strongest objection, make the causal link visible, remove claims that create needless resistance and leave enough space for the listener to complete the reasoning. The aim is not to occupy the mind. It is to give it better material to work with.

Feeling and Belonging Set the Stakes

Logic and emotion are often presented as rival departments, while social influence is treated as a separate chapter about crowds. Human judgement does not respect those divisions. Emotion marks what matters. Other people help decide what is safe to believe, costly to say and possible to do.

Different emotions organise attention and action differently. Fear highlights danger and escape. Anger assigns blame and can support approach rather than retreat. Guilt concerns conduct and can motivate repair. Shame threatens the whole self and often produces hiding or defence. Hope makes delayed benefit feel available. Disgust marks contamination and can spread from an act to a person or group. These are tendencies, not fixed programmes. People, cultures and settings differ.

Fear appeals show why technique without mechanism misleads. A 2015 meta-analysis covering 127 papers and more than 27,000 participants found a positive average effect on attitudes, intentions and behaviour. Effects were stronger when messages included information about an effective response. The result supports neither popular extreme. Fear does not usually doom a message, and more fear is not always better. Fear says there may be a fire. Efficacy shows a usable door.

The social environment supplies another kind of evidence. A full restaurant, a widely adopted practice or expert agreement may reveal information an individual lacks. But a crowd can be impressive without being informative. If everybody copied the first person, its size exaggerates the amount of independent judgement behind it. Selection by an algorithm or display by the seller gives further reason to inspect the clue. Such signals may still

persuade. Their psychological force and their evidential value are different questions.

Solomon Asch made the pressure visible with a task that should have needed little persuasion: matching line lengths. Other people in the room, secretly working with the researcher, unanimously gave a wrong answer. Many responses remained independent and correct. Some participants who yielded doubted their own judgement; others went along despite their misgivings. The interviews revealed more than one route to conformity. A wrong public answer could reflect uncertainty, the discomfort of standing apart, or both. Agreement told the observer less about private conviction than it appeared to.

Descriptive and injunctive norms separate prevalence from approval. Telling people that a harmful behaviour is common may intend condemnation while advertising permission. A field experiment gave California households feedback comparing their electricity use with neighbours. High users tended to reduce consumption, while some low users moved upwards towards the norm. A small approval cue for low consumption helped prevent that boomerang. The information said both *this is common* and *this is valued*. The second message changed the direction of the first.

Identity joins feeling to the social price of revision. A position can signal loyalty, competence, morality or adulthood. Accepting a factual correction may then threaten more than a proposition. Humiliation raises the cost further. A message that wins the evidential point while making agreement look like betrayal can lose the person. Private space, a credible in-group messenger or an honourable route to revision may matter more than extra force.

Groups can misread themselves. People may privately doubt a norm while assuming everybody else accepts it. Public silence then manufactures consensus. Revealing representative private views can change conduct without inventing a new argument, though unrepresentative numbers merely replace one false social fact with another.

Emotion and belonging also determine what survives. A brief emotional surge may move an immediate choice and vanish. A new practice is more likely to persist when later encounters make it normal, feasible and compatible with membership. Networks can reinforce a tentative update or return the person to the old answer. Persuasion at scale therefore depends on institutions, peers and environments as much as on the message that began it.

Good persuasion makes stakes visible without manufacturing them. It matches emotional force to evidence, separates what people do from what they approve, and asks what the recipient would risk by agreeing. When feeling or social pressure is doing more work than the claim can support, conviction has been borrowed rather than earned.

Ownership and Experience Decide What Lasts

Persuasion asks a person to revise while preserving the sense that the mind remains theirs. Mishandle that condition and the attempt creates its own opposition.

Psychological reactance begins when a freedom feels threatened. The response can combine anger with counterarguing and a motive to restore the threatened option. Commands, controlling language, censorship and conspicuous pressure can therefore make the rejected choice more attractive or the source less credible. The effect varies with the importance of the freedom, the person and the setting. It is not a law that every firm message backfires. It explains why increasing force can reduce influence when autonomy, rather than evidence, is the blocked stage.

This gives resistance a positive role. A mind that never resisted would accept fraud, gossip and bad advice with equal ease. Suspicion, source checking and counterargument are part of the immune system of belief. Persuasion succeeds more honourably when it survives those defences than when it slips past them.

Inoculation theory uses that idea directly. Expose a person to a weakened challenge, warn that the position may be attacked and let them rehearse

replies. A meta-analysis of inoculation research found that this can increase resistance to later counter-persuasion across the studied settings. The medical metaphor has limits: messages are not viruses, and protection may decay or fail to transfer. The durable lesson is that rebuttal skill can be taught before the pressure arrives.

Agency also explains self-persuasion. People are often more convinced by reasons they generate, choices they enact and conclusions they can attribute to themselves. Cognitive dissonance research examines the discomfort created when beliefs, self-image and conduct conflict. In the famous 1959 forced-compliance experiment, people paid a small sum to describe a dull task favourably later rated it more positively than those paid enough to justify the lie as a job. The study was small and its exact mechanism remains debated, but it helped redirect persuasion science towards the recipient's effort to make sense of their own behaviour.

Action can change attitude rather than merely express it. Trying a product, attending one meeting, signing a pledge or helping a cause supplies experience and a new question: why did I do that? This can be benign when the first step is transparent and freely chosen. It becomes manipulation when a trivial commitment is engineered to conceal an escalating demand. The same mechanism that can support a healthy routine can also recruit consistency to defend a trap.

Conversation can create ownership by adapting to objections and letting reasons emerge. Structured canvassing has changed selected exclusionary attitudes, but those results belong to complete encounters, not to listening alone. A person can feel understood, like the speaker more and remain unconvinced. Listening still has a practical use: it can reveal whether the obstacle is an implausible prediction, a feared cost or an objection the speaker has missed. That information helps only if the answer changes what happens next. Repeating the original pitch more warmly is still repeating the original pitch.

Ownership is still insufficient if action is impossible. An intention must survive cost, time, skill, defaults, reminders and the availability of the

required object. A person can be convinced that recycling matters and still lack a food-waste caddy. They can accept medical advice and be unable to afford the journey. Removing friction is not persuasion in the narrow sense, but a message that ignores friction mistakes agreement for completion.

Experience then judges the claim. A trial can make an abstract benefit visible, but it can also expose exaggeration. A new action may strengthen an attitude when the result is rewarding and intelligible. It may weaken it when the promised consequence fails to appear. Persistence depends on what happens after the persuasive encounter, including feedback, habit, social support and whether the original source remains credible.

This closes the loop. A persuasive attempt begins with prior beliefs, trust, identity and experience. What follows changes those same starting conditions. A fulfilled promise gives the next message an easier hearing. A hidden cost teaches suspicion. A humiliating victory hardens resistance. A freely owned change can become part of the person's evidence and self-conception.

Ethical persuasion leaves the recipient able to state the case, inspect its source and refuse it. Pressure may produce compliance while the threat is present. Ownership, feasible action and confirming experience give change reasons to persist after it leaves, though none guarantees permanence. What makes a mind hard to change can also make a change last: people defend what feels theirs.

How It Actually Works

The case before the science

Aristotle treated persuasion as a practical problem of judgement. A speaker faces an audience that must decide under incomplete information, limited time and consequences that cannot be settled by proof alone. The task of rhetoric was to identify the available means of persuasion in a particular case. The last phrase matters. There was no master key, because each case came with a different audience, question and standard of decision.

His familiar triad still supplies a useful map. Logos concerns the case being made. Ethos concerns what the audience infers about the speaker's character and competence. Pathos concerns the audience's emotional condition. These are not three ingredients ladled into a speech. They interact. A concession can strengthen the logic and the speaker's credibility. A frightening fact can alter both feeling and the amount of scrutiny. A trusted speaker can make an uncertain claim worth examining without making it true.

Aristotle also centred the enthymeme, an argument whose force depends partly on premises the audience already accepts. If a court agrees that a person who fled had something to hide, the speaker may not need to state that premise. Modern textbooks sometimes call the enthymeme an incomplete syllogism. In practice it is an audience-linked argument. Its missing material is supplied by shared knowledge, values or assumptions. That gives rhetoric its efficiency and its danger. What the audience supplies may be false.

For centuries, persuasion remained chiefly a craft studied through speeches, legal argument, sermons and written models. The audience was often visible. An orator could watch attention fade, hear opposition and adapt. Printing, newspapers, film and broadcasting changed the geometry. One message could reach people the sender could not see, at a scale where direct response was impossible. The audience became an estimate, and the effect became a question that eloquence alone could not answer.

The films that taught facts

The Second World War turned that question into an organised research programme. The United States Army commissioned the *Why We Fight* films to explain the conflict to soldiers and strengthen commitment. Carl Hovland and colleagues in the Army's Research Branch tested parts of the programme rather than assuming that patriotic purpose guaranteed persuasive success.

The broad finding was awkward for anyone who believed information

naturally becomes motivation. The films could improve factual knowledge. Their effects on broader attitudes and willingness to serve were smaller, uneven or absent. A soldier could learn why the war had begun without feeling differently about the Army, the enemy or the length of service ahead. The outcome ladder had appeared in data: reception, learning and attitude change were separable.

After the war, Hovland's Yale group broke communication into questions that could be varied and measured. Who is speaking? What are they saying? Through which arrangement? To whom? With what result? Researchers compared more and less credible sources, one-sided and two-sided messages, explicit and implicit conclusions, fear levels, order effects and audience differences. The programme made persuasion experimental.

Its strength was control. Change one feature, hold others steady and compare groups. Its weakness was the same. A brief message in a study can isolate a mechanism while stripping away reputation, history, social consequences and repeated exposure. The clean result may describe what happens under those conditions and little else. Persuasion science has been negotiating that trade ever since.

Early models often resembled a relay. The recipient must attend, comprehend, accept and retain. Each stage is necessary for some later outcomes, so the sequence remains useful. Yet it can make the listener look passive, as though a message moves through mental checkpoints until belief is deposited. The next shift in the field came from noticing that recipients do something while the message is passing through them. They answer back.

The thoughts between message and judgement

A listener does not store an argument whole. They connect it to memories, test examples, notice omissions, infer motives and generate replies. These cognitive responses can matter more than the words as counted on the page. A strong message creates favourable thoughts that survive inspection. A weak message can become less persuasive when attention rises because careful thought reveals the defects.

This is why distraction has no fixed effect. If a message is strong, distraction can prevent the recipient from appreciating it. If the message is weak and the audience would otherwise counterargue, mild distraction may temporarily protect it. The same variable changes direction because persuasion depends on what thought it interrupts.

In a 1981 experiment, American undergraduates heard a case for compulsory comprehensive exams before graduation. For some, the proposal would take effect the following year. For others, it was ten years away: somebody else's problem. Researchers varied the strength of the arguments and the expertise attributed to their source. When the proposal was presented as affecting the students themselves, argument quality mattered more. When it lay safely beyond their university years, the source's credentials carried more weight. The issue had not changed. Its personal consequences had.

The elaboration likelihood model developed by Richard Petty and John Cacioppo explains this shift in terms of motivation and ability to think. Shelly Chaiken's heuristic-systematic model makes a related distinction between examining a case and using economical rules of judgement. Neither assigns people permanently to a logical or gullible camp. The student impressed by credentials today can become the student demanding better evidence tomorrow. Nor must a cue be mindless: the expertise of a source can be a sound reason to investigate a claim.

The mode of processing can affect what survives. Attentively examined attitudes are often more stable than impressions resting on a cue that later

disappears. But effort is no warranty: lengthy thought can entrench error when someone selects evidence to defend a preferred answer.

The practical question becomes: what mental work does this message invite? A comparison table permits checking. A worked example lets the listener follow a consequence through. Neither helps much if the person has no reason, or no opportunity, to use it.

How principles became buttons

As experimental findings accumulated, popular persuasion compressed them into principles. Robert Cialdini's influential framework describes recurring forms of influence: reciprocity, commitment and consistency, social proof, authority, liking, scarcity and, in the expanded edition, unity. Each identifies a real source of information or social pressure. Trouble begins when a conditional tendency is advertised as a button.

Reciprocity helps social exchange work. A gift, favour or concession can create an obligation to answer in kind. Its effect depends on whether the first act looks voluntary, proportionate and relevant. An unwanted gift attached to a demand can produce irritation rather than gratitude. Consistency lets people maintain a coherent identity and honour commitments, yet a commitment imposed publicly or obtained by concealment is more likely to produce performance, resentment or escalation than considered belief.

Authority economises on expertise: a patient cannot repeat the training behind every recommendation. Relevant competence is informative; a uniform or confident tone can imitate it. Scarcity may signal demand or a closing opportunity, but can also be staged. Liking can encourage a hearing, while obvious flattery exposes the attempt. Social proof draws force from what others appear to believe or do, including appearances that turn out to be manufactured. Unity goes beyond liking: someone perceived as *one of us* may speak with the authority of shared identity. That makes belonging influential, not infallible.

These principles often influence compliance, selection or willingness to continue. That is important, but it is narrower than changing a belief. A

customer may accept a free trial because the default is easy, buy before a deadline and regret it after the pressure lifts. A person may comply with an authority while privately disagreeing. The technique changed action under a condition. It did not necessarily create an attitude that will travel.

Use the principles to inspect an encounter rather than operate a control panel. A favour may explain an urge to reciprocate; it does not settle whether the request is reasonable. A deadline may explain urgency; it does not prove the opportunity will disappear. Recognising the source of pressure creates room to examine the case beneath it.

The field refuses the laboratory average

Outside the laboratory, persuasive attempts compete with prior contact, opposing messages, selective exposure and events. The sender may fail to reach the intended person. The person may skim, distrust, forget or encounter a rebuttal five minutes later. The outcome may depend on whether friends approve, whether the requested behaviour is possible and whether the next experience confirms the claim.

Political campaigns make these problems visible because they spend heavily, measure often and operate in hostile information environments. A 2018 meta-analysis of field experiments found that campaign contact in United States general elections had an average effect on candidate choice close to zero in the settings studied. That result is frequently inflated into the claim that political persuasion does not work. It supports something narrower: late campaign contact aimed at an already saturated choice is a hard setting in which average effects can be tiny even when particular messages or subgroups move.

Other field experiments have found durable changes on selected exclusionary attitudes after structured, two-way conversations. Those interventions combined several elements: a human messenger, personal narrative, questions, sustained attention and encouragement to reflect on experience. Their success does not reveal one magic ingredient. It shows that a complete encounter can sometimes move a socially charged attitude

where a broadcast appeal would struggle.

A preregistered field experiment published in 2025 tested listening more directly. Researchers analysed 1,485 United States participants' responses after video conversations about in-state tuition for undocumented immigrants. Adding high-quality, non-judgemental listening to a persuasive narrative improved liking, trust and feeling understood, and reduced defensiveness. It produced no detectable additional attitude change, either immediately or five weeks later. The appeal was not tailored to what the listening revealed, so this was not a test of every possible use of conversation. It did separate a warmer relationship from a more persuasive argument. A humane practice does not need a false effect to justify it.

A persuasive encounter is therefore more than the sum of its attractive features. Listening may help one outcome, a narrative another. Testing the combination does not reveal which ingredient caused which change. Testing an ingredient does not reproduce every way it could be used. The useful question is precise: what did this addition improve, beyond the encounter it was compared with?

Changing an existing false belief

Correction is persuasion with an occupied target. The recipient has a claim, an explanation and perhaps a social reason for keeping both. Removing one sentence leaves a gap. If the original account remains easier to retrieve and explains the event, it can continue to influence judgement even after the person accepts that it was false.

Research on misinformation rejects two pieces of folklore at once. Corrections often reduce false belief. Reliable backfire, in which correction routinely strengthens the falsehood, is much rarer than popular discussion suggests. Yet correction is not deletion. Residual influence, repetition, source distrust and identity can limit the effect, and a later memory may preserve the claim while losing the warning attached to it.

An alternative explanation gives the mind something to use instead. Imagine a report blaming a train cancellation on a driver's absence. Saying the driver

was present removes an accusation but leaves the cancellation unexplained. Establishing that a signal failure stopped the service supplies a replacement cause. The example is hypothetical; the principle is practical. A correction should make the accurate account clear enough to answer the question the false account had answered, even badly.

Timing matters. Inoculation can warn people about a misleading tactic before exposure and let them practise resistance. A correction after exposure can still work, especially when it is prompt and explanatory. Repetition may be needed because the false claim has already received repeated rehearsal. The ethical aim is not to make the correction equally loud. It is to make the accurate model easier to retrieve when the next judgement occurs.

Correction also exposes the difference between defeat and revision. Humiliating somebody for having believed a false claim raises the identity cost of accepting the correction. Pretending that all interpretations are equally sound abandons the truth. The harder route is to make the evidential standard firm and the social route back open.

Adaptive persuasion

Digital communication made experimentation cheap. A sender can vary a headline, image, order or appeal, observe clicks and keep the version that produces more of them. That system can optimise a measured response without understanding why it changed. It can also optimise the wrong rung. A message that wins clicks through alarm may reduce trust, increase cancellations or spread a false inference.

Personalisation adds another layer. Matching a message to expressed values, circumstances or prior behaviour can increase relevance. It can also feel invasive, divide audiences into inconsistent realities and make the sender's private knowledge part of the persuasion. The ethical question is not settled by whether the match improves an average effect. It includes how the data were obtained, what is concealed, whether important information is withheld from one group and whether the recipient can

recognise the attempt.

Generative systems reduce the cost of producing and adapting language. Controlled studies now show that they can move stated policy attitudes under defined conditions, while also showing why one headline cannot summarise the field. In a debate study with 900 United States participants, GPT-4 given basic personal information outperformed the human-human baseline. An author correction published in September 2026 clarified that the direct comparison with unpersonalised GPT-4 did not reach statistical significance. The experiment therefore supports the capability of the personalised system against the human baseline, but does not isolate personalisation as the cause of its advantage.

A larger 2025 programme tested 19 models with 76,977 participants across 707 political issues. Post-training and prompting increased persuasiveness more than greater model scale or personalisation in those experiments. The same interventions that raised persuasion also reduced factual accuracy among the claims checked. Optimising for persuasive success can therefore reward less accurate content. The metric can improve while the message gets worse.

A 2026 study separated persuasive capacity from reach. In two survey experiments with 10,417 participants, chatbot conversations were about as persuasive as campaign adverts among those exposed. The authors then simulated campaigning at scale, accounting for whether people would encounter and engage with the messages. Under those assumptions, traditional methods reached more people effectively. That was a modelling result, not a measured contest between deployed campaigns. A system may be persuasive in conversation and still struggle to get the conversation started.

These findings concern selected political topics, stated attitudes and particular systems, prompts and comparisons. They do not establish universal superiority, durable behaviour change or effortless mass manipulation. The sender can run many attempts, retain the winners and hide the failures. Scale changes risk even when the average individual

effect remains modest, especially when optimisation rewards agreement without truth.

The operating model remains human. Define the outcome. Diagnose the starting position. Secure attention and comprehension. Establish a source relationship. Prompt useful thought and proportionate feeling. Account for social consequences. Make the action feasible. Preserve enough agency for the recipient to own the conclusion. Then test the outcome after the encounter has ended.

The sequence in one encounter

As an illustration, imagine a council wants households to separate food waste. The desired outcome is not approval of recycling, recall of a leaflet or an intention reported at the door. It is repeated sorting on collection days. That definition changes the persuasive job.

The starting audience may include committed recyclers, indifferent residents, tenants with little kitchen space and people who distrust the council. One universal appeal wastes those differences. The message first has to arrive before the caddy does, explain what belongs inside and make the schedule easy to retrieve. Confusing instructions can defeat an audience that already agrees.

Credibility depends on conduct as well as copy. Residents need to believe the waste will be processed as described and that reported costs are not decorative. A clear explanation of what happens after collection can supply a causal model. Environmental benefit may matter to one household, cleaner bins or lower disposal costs to another. None requires a false claim.

Social information must be chosen carefully. Saying that most neighbours participate can reduce uncertainty if the figure is accurate. Advertising widespread non-compliance would teach the opposite norm. A respected caretaker or local resident may be a better messenger than a distant official, though no messenger can make an inconvenient system convenient. The caddy, liners, storage, collection frequency and reminders determine whether intention survives Tuesday morning.

Agency remains visible. Residents can ask questions, report problems and understand why the rule exists. A threat may secure one week's compliance while generating evasion later. A voluntary commitment may help when the practical path is easy and the choice is honest.

Evaluation then follows the outcome. Compare similar areas with and without the intervention, measure contamination and sustained participation rather than survey warmth, and look for costs displaced elsewhere. The example is hypothetical. Its purpose is to show that persuasion is distributed across message, source, setting, action and feedback. A brilliant slogan cannot repair a missing bin.

How we know

Persuasion research combines rhetoric, laboratory experiments, surveys, field trials, meta-analyses and observations of behaviour. Each reveals a different part of the process. Experiments identify causal effects under controlled conditions. Field trials test interventions amid competing messages and practical friction. Longitudinal designs show whether change lasts. Meta-analyses estimate average patterns while inheriting variation in measures, populations and publication quality.

The evidence base is uneven. Many classic studies used Western students, brief exposures and self-reported attitudes. Effects on one policy, product, health choice or identity conflict do not transfer automatically to another. Belief, attitude, intention, compliance and behaviour are often measured as though interchangeable. They are not. Small average effects can matter at scale, while impressive immediate shifts can disappear before any action occurs.

Current claims about conversational artificial intelligence are especially provisional because models, deployment and exposure change faster than long-term evidence can accumulate. One prominent 2024 study of AI conversations and conspiracy belief received an editorial expression of concern in 2026 and is not used here as support. The strongest conclusion is therefore conditional: mechanisms recur, but their size and durability must

be tested in the setting where they will be used.

What People Get Wrong

“Facts do not change minds”

This slogan usually begins with a true experience. Someone received accurate evidence and remained unmoved. The conclusion then outruns the event: information is powerless, belief is tribal and only stories or identity can reach people.

Facts change minds often enough for education, diagnosis, journalism and ordinary correction to exist. Experiments generally find that clear corrections reduce belief in false claims. New evidence can alter estimates, attitudes and choices. The effect is weaker when the source is distrusted, the evidence seems irrelevant, the claim protects belonging or the requested behaviour is difficult. None of that makes the factual content optional.

The slogan confuses insufficiency with uselessness. A laboratory result does not prescribe its own meaning. A risk figure needs a baseline. An expert claim needs a reason to trust the expert. A correction may need an alternative explanation. An accurate proposition can fail because another part of the persuasive chain is blocked.

Fatalism has a cost. If facts never matter, communicators are licensed to replace evidence with emotional pressure and to treat disagreement as pathology. The better rule is harder: tell the truth, then diagnose what else must be true for the truth to be understood, trusted and used.

“You persuade with logic or emotion”

The opposition sounds clean because arguments contain reasons and feelings feel different from reasons. Human judgement does not keep them in separate rooms.

Emotion marks significance. Fear says a threat deserves attention. Anger assigns blame. Hope makes effort appear worthwhile. Compassion enlarges

whose welfare enters the decision. These appraisals can be accurate or distorted, just as a calculation can use sound or selected inputs. Logic organises relations among claims; it cannot decide which consequences a person should care about without some account of value.

Fear-appeal research makes the interaction visible. Threatening messages have produced a positive average persuasive effect across studied settings, especially when recipients are also given an effective, feasible response. That does not mean warnings without response information normally backfire; the meta-analysis did not find that reversal. The feeling is neither the proof nor the enemy of proof. It changes what the proof means for action.

The useful distinction is between warranted and unwarranted influence. Does the emotional force fit the evidence? Does it direct attention towards material consequences or away from missing information? Does the reasoning survive after arousal falls? Strong persuasion joins an inspectable case to stakes the audience can feel. Weak persuasion asks one to disguise the absence of the other.

“Body language carries 93 per cent of meaning”

The formula gives 7 per cent to words, 38 per cent to tone and 55 per cent to facial or bodily cues. It is memorable, precise and wrong as a general law.

The numbers were derived by combining two small studies published by Albert Mehrabian and colleagues in 1967. They concerned judgements of liking or attitude when spoken words, vocal tone and facial information were inconsistent. Participants were not learning a legal rule, following technical instructions or deciding whether an argument was valid. Mehrabian did not establish that words carry only 7 per cent of communication.

Non-verbal behaviour still matters. Tone can reveal uncertainty, contempt or warmth. Facial expression can alter how an ambiguous sentence is read. Mismatch between a person's words and manner can reduce trust. Yet the content determines whether a medicine dose is two milligrams or twenty, whether a contract renews and what evidence supports a claim. No posture

can supply the missing zero.

The myth survives because training courses prefer a pie chart to a boundary condition. Its correction matters beyond trivia. It stops communicators spending effort on performance while neglecting the proposition, and it stops observers pretending that an impression of confidence is a measurement of truth.

“Subliminal messages can make you do anything”

The founding story was a fraud. In 1957, market researcher James Vicary claimed that imperceptible cinema messages telling viewers to buy popcorn and Coca-Cola had sharply increased sales. He produced no adequate record, could not reproduce the result and later admitted that the supposed experiment was a promotional gimmick.

Subliminal perception is not therefore imaginary. Stimuli presented below conscious report can prime limited responses under controlled conditions. A 2006 experiment found that subliminal exposure to an iced-tea brand affected later choice chiefly among participants who were thirsty. The boundary is the finding: a relevant cue interacted with an active goal in a constrained choice. It did not implant a new thirst, loyalty or purchase against the person's interests.

Claims for subliminal self-help have fared badly. In a double-blind study of commercial tapes, labels shaped what users thought had improved, while the hidden content did not deliver its promised memory or self-esteem benefits.

The fantasy persists because invisible control explains choices without the discomfort of ordinary influence. In practice, visible systems are more capable: repetition, defaults, selective information, social pressure and personalised appeals. Looking for a secret frame can distract from the message and environment operating in full view.

“Correcting misinformation usually backfires”

The backfire idea became popular because it fits a dramatic model of identity defence: challenge a false belief and the believer grips it harder. Such reactions can occur. They are not the standard result of correction.

Reviews and large experiments find that corrections generally move belief towards accuracy. Researchers have struggled to reproduce a broad, reliable backfire effect across political claims and audiences. That does not mean correction is complete or permanent. The original claim can continue to influence reasoning, especially when it supplied a causal explanation. Repetition can preserve familiarity. Distrust can limit acceptance. A correction may fade faster than the rumour.

Poor correction also creates avoidable problems. Repeating the falsehood as the headline can make it fluent. Ridicule raises the social cost of revision. A bare denial leaves the explanatory gap open. Better corrections foreground the accurate account, explain the error, provide a replacement model and use a source the audience can examine.

The myth matters because communicators sometimes withhold corrections for fear of making things worse. Silence grants the false claim uncontested rehearsal. Caution should improve the correction, not abolish it.

“One clever principle works on anyone”

Reciprocity, scarcity, authority, social proof, liking and commitment are often presented as switches. Press the right one and a predictable person complies.

The first distinction is between a clue being persuasive and being sound. Independent expert agreement can deserve weight. A fake review count may also influence someone who has not recognised it as fake. The same goes for a real deadline and an invented one. These mechanisms explain why an appeal might work; they do not certify its honesty. Nor do they ensure success: recipients bring different knowledge, motives and suspicions to the encounter.

Change the conditions and the direction can change. A countdown timer caught resetting damages the claim of scarcity and the seller's credibility. Manufactured popularity produces suspicion when discovered. Excessive flattery exposes the speaker's motive. A small public commitment can support follow-through or trap someone in escalation.

The principles became buttons because buttons are easy to teach and sell. Their scientific value lies elsewhere: they identify questions about information, obligation, identity and uncertainty. The person applying a universal script has stopped looking at the audience, which is the quickest way to misuse every item on the list.

“If they agreed, you persuaded them”

Agreement is an event. Persuasion is a claim about what changed.

A person may agree to end an argument, please a superior, avoid embarrassment or obtain a reward. They may understand and believe a factual statement while retaining the same evaluation. They may form an intention that loses to habit, cost or inconvenience. They may act once under observation and revert when the observer leaves. Each outcome is real. None proves the others.

This is why immediate self-report is a weak final measure. It can capture a genuine first step, but it also captures politeness, demand characteristics and temporary arousal. Stronger evidence matches the promised outcome: later recall for learning, private attitude when public conformity is possible, observed choice for behaviour, repeated action for persistence and a comparison group for causation.

The distinction also protects ethics. Coercion can produce behaviour. Manipulation can produce a signature. Neither should be renamed persuasion merely because the target moved. The most convincing agreement is the one that survives privacy, delay, practical friction and the removal of pressure. Until then, the speaker may have won the room and nothing beyond it.

Use It

Name the change you are asking for

Before choosing words, specify the outcome. Do you need the person to notice, understand, believe, approve, intend, act once or continue? These are different jobs.

A safety notice may need accurate recall. A medical discussion may need informed choice rather than agreement. A public campaign may aim to correct a false belief, while the desired behaviour still depends on cost and access. A manager may need staff to use a new process after the meeting, not say that the presentation was clear. The outcome determines the evidence of success.

Write the requested change as a sentence with a time and an observable test. "People will find this persuasive" is circular. "Recipients will identify the correct dose tomorrow" or "more eligible households will complete the application within a month" can be examined.

This discipline prevents the most common inflation in the subject. Attention is reported as interest, interest as intention and intention as behaviour. Name the rung and you stop claiming the ladder.

Find the blocked stage

When persuasion fails, communicators often add force to the message they already prefer. More statistics, a stronger warning, a longer speech. Diagnosis comes first.

If the audience never encountered the claim, the problem is reach. If they cannot repeat it accurately, the problem may be comprehension. If they understand and reject it, inspect evidence, source and prior belief. If they agree privately but not publicly, look at social cost. If they intend but do not act, look at feasibility, timing and habit. If change fades, look at later experience and reinforcement.

Ask questions that separate these possibilities. What do you think I am

proposing? Which part seems doubtful? What would accepting it imply? What evidence would count? What makes the action hard? The answers may expose a disagreement over facts, values, trust, identity or implementation. Do not treat all five as ignorance.

A blocked-stage analysis also tells you when to stop talking. If the obstacle is an unaffordable price, a missing appointment or a system that breaks its promise, another paragraph is camouflage. Repair the condition or narrow the claim.

Make credibility inspectable

Do not ask for trust as a substitute for the conditions that create it.

State your role, interest and uncertainty. Distinguish what the evidence establishes from what you infer or recommend. Use sources the recipient can reach. Explain why the source has relevant competence and what would expose an error. Where possible, make a prediction specific enough to check later. An admission that weakens your immediate case can strengthen the relationship if it shows that accuracy outranks victory.

Match messenger to domain. A respected peer may know how a change affects daily work. A technical specialist may know whether it functions. A leader may hold the decision right. None automatically possesses the others' knowledge. Combining them can be more credible than dressing one person in universal authority.

Watch for borrowed signals. Confidence, fluency, status, professional design and numerical precision can all increase credibility without increasing truth. Ask what process produced the claim, what comparison was made and what incentives surround it.

Trust built this way is slower than a slogan and more portable. It can survive a difficult message because the recipient has grounds for judging the source rather than an instruction to believe it.

Give the recipient useful work

Persuasion is stronger when the audience performs part of the reasoning. That does not mean hiding the conclusion or setting a conversational trap.

Invite comparison. Ask the person to forecast what each option would produce, retrieve a relevant experience or explain the strongest objection. Present enough evidence for them to test the claim. Questions reveal the thoughts occurring between message and judgement, which lets you answer the real counterargument rather than the convenient one.

Make the opposing case recognisable. A token objection followed by an easy demolition advertises that the process is staged. State the serious cost, uncertainty or rival explanation, then show why you still reach your conclusion. A two-sided message can increase credibility when the audience knows an objection exists, but it fails if the reply is weak or the concession hides a larger problem.

Leave room for refusal and revision. Controlling language can turn autonomy into the issue. Offer a clear recommendation where responsibility requires one, then distinguish recommendation from compulsion. Ask the recipient to put the case in their own words and identify what would change their mind later. Ownership is stronger when the person can explain both the reason and its limits.

Lower the price of changing sides

A person can accept the evidence and reject the humiliation attached to accepting it.

Correct privately when public reversal would threaten status. Give credit for the concern that made the old position reasonable. Separate “this claim is false” from “you were foolish to believe it”. Use messengers who show that revision is compatible with membership, especially where a belief has become a group marker. Do not demand a confession as the entrance fee to accuracy.

This is not permission to blur truth. Face-saving becomes dishonest when it

pretends that evidence is equal or conceals responsibility for harm. The task is to remove needless social punishment, not factual consequence. A person can acknowledge damage and still be offered a route forward.

Commitment should receive the same scrutiny. Public pledges can support action, yet they can also make reconsideration costly. Keep commitments specific, informed and reversible where the decision warrants it. Praise correction as consistency with a deeper standard, such as accuracy or care, rather than as betrayal of a previous self.

The best sign of a healthy persuasive setting is that people can update without becoming socially homeless.

Measure the afterlife

The effect that matters often occurs after the message ends. Build the bridge from judgement to action, then look later.

Make the requested behaviour feasible. Specify what happens, when and where. Remove avoidable friction. Supply the tool, form, appointment or reminder that an intention needs. This is where persuasion meets environment: a convinced person can still fail inside a badly designed system.

Compare outcomes against a credible counterfactual. A before-and-after rise may reflect selection, season or another event. Random assignment is strongest where practical and ethical; matched comparisons and interrupted trends can still improve on applause. Record unintended effects. A norm message may improve high users and worsen low users. A frightening appeal may increase stated concern and avoidance.

Measure more than the easiest response. Clicks can reward alarm. Signatures can reward pressure. Immediate agreement can reward politeness. Use delayed belief, observed behaviour, repetition, cancellation, trust and error correction where the promise requires them. Report the denominator and the people who never received the message.

Persuasion is judged by the change claimed, not the metric available.

The limits

Persuasion cannot settle every disagreement. Evidence may be incomplete. Values can conflict without one side failing to understand. Interests may oppose each other. A person may rationally reject a proposal that benefits the speaker. Respect for agency includes the possibility that the answer remains no.

The same mechanisms can improve falsehood. Clear language, trusted sources, identity, repetition and social support do not belong to truth. A skilled communicator can make a bad case easier to process. Ethical persuasion therefore requires accurate claims, material disclosure and a real option to refuse. Manipulation and coercion may move behaviour more strongly; that does not turn them into superior persuasion.

Effects are uneven. A result from one population, issue or channel is not a universal coefficient. Conversation can be impossible at scale. Time can be scarce. Some decisions require rules, enforcement or institutional design rather than another attempt to win individual assent. Communication should not be used to decorate an unjust choice or transfer responsibility to the people carrying its cost.

The model improves diagnosis. It does not grant control.

The one thing to keep

Keep the inside view.

Imagine a colleague rejecting a proposal you have spent weeks preparing. The tempting verdict is that they do not understand it. But perhaps they understand the unpaid work it creates better than you do. Finding that out changes the task: the proposal needs repair, not a more forceful presentation. Or their objection may rest on a factual error that evidence can correct. The inside view is the discipline of discovering which problem you have before deciding how to overcome it.

This gives you a reason to listen that no promise of influence can supply. You might learn that the other person has a good reason to refuse. You

might find a better case, or discover that your own position should change. The recipient stops being a lock and becomes an evaluator; so do you. A worthwhile argument leaves both minds available for revision.

Minds are changed from the inside. Persuasion can arrange the conditions. The other person still does the changing.

Terms

A glossary of the distinctions that keep persuasion from collapsing into a bag of tricks.

Persuasion. An intentional attempt to alter another person's belief, attitude, intention or voluntary behaviour through communication. It differs from coercion because refusal remains a meaningful option. It can succeed without producing immediate agreement.

Belief. A judgement about what is true or likely, such as whether a treatment works. Belief can change without changing whether the person approves or acts. Beliefs also differ in confidence and accessibility.

Attitude. An evaluation of an object, person, policy or behaviour, ranging from favourable to unfavourable. Attitudes vary in certainty, importance, accessibility and connection to identity. An attitude may be ambivalent rather than uniform.

Intention. A stated or formed readiness to perform a behaviour. Intentions predict action imperfectly because opportunity, habit, competing goals and practical friction intervene. Strong intentions are still plans, not completed acts.

Behaviour. What a person does rather than reports, remembers or intends. Persuasion studies weaken their claims when an attitude rating is treated as proof of action. Repeated behaviour is a stricter test than one response.

Compliance. Agreeing to a request or behaving as asked, often in a specific situation. Compliance may occur without private belief and may vanish when pressure disappears. A reward, threat or authority can secure

it.

Attitude strength. The collection of qualities that make an attitude durable and consequential, including certainty, importance, accessibility, knowledge and resistance. Strong attitudes more often guide behaviour. Direct experience often increases some of these qualities.

Rhetoric. The disciplined study and practice of persuasion in situations requiring judgement. In Aristotle's account, it identifies the available persuasive means in each particular case. It concerns probable decisions, not mathematical proof.

Ethos. Persuasive force arising from the character, competence and intentions an audience attributes to a speaker. Ethos is granted by an audience, not owned permanently.

Pathos. Persuasive force arising from the audience's emotional condition. Emotion alters attention, value and perceived risk, but it does not establish whether a claim is true.

Logos. The case supplied by reasons, evidence and relations among claims. Logos can include examples and probabilities as well as formal deduction.

Enthymeme. A rhetorical argument that depends on premises the audience supplies from shared knowledge or values. Its economy is powerful; its hidden assumptions may also be false.

Source credibility. The recipient's judgement that a communicator knows enough and can be trusted. Relevant expertise, honesty, goodwill and accountability can contribute in different settings.

Elaboration. The amount and direction of issue-relevant thought a person gives a message. Greater elaboration helps strong arguments and can expose weak ones.

Elaboration likelihood model. A theory explaining how motivation and ability shape message processing. It contrasts more effortful examination

with greater reliance on accessible cues, without proposing literal brain routes.

Heuristic. A simplifying rule such as trusting experts or following a majority. Heuristics conserve effort and can be sensible, but only when the cue is informative.

Framing. Selecting and organising true features of an issue so that some interpretation becomes salient. Framing becomes deceptive when it hides material information or changes the comparison unfairly.

Cognitive response. A thought generated while receiving a message, including support, objection, inference or memory. These responses often mediate between the words heard and the judgement formed.

Motivated reasoning. Processing shaped partly by a desired conclusion, identity or emotional need. The term does not imply that accuracy motives disappear or that all disagreement is defensive.

Identity-protective cognition. Reasoning that helps preserve membership or self-conception when evidence threatens a group-linked belief. Revision then carries social and psychological costs beyond the factual claim.

Social proof. Using other people's choices as evidence about what is correct or valuable. Independent, informed choices are better evidence; copied or manufactured popularity can still exert influence.

Descriptive norm. A belief about what people commonly do. Publicising an undesirable descriptive norm can normalise it, even when the communicator intends to condemn it.

Injunctive norm. A belief about what people approve or disapprove. It concerns social evaluation rather than prevalence and can counteract a misleading descriptive norm.

Conformity. Adjusting judgement or public response towards a group. Conformity may reflect accepted information, avoidance of social cost or both, so public agreement need not reveal private belief.

Reactance. Anger, counterargument and freedom-restoring motivation triggered when a choice feels threatened. Controlling language can therefore create opposition beyond the substance of the request.

Inoculation. Building resistance by warning of an attempted challenge, presenting a weakened version and rehearsing replies. Protection varies by message, delay and transfer to new attacks.

Mere exposure. Increased familiarity, and sometimes liking, produced by repeated contact with a stimulus. Familiarity can help an unknown option, but repetition can also bore or irritate.

Illusory truth effect. The tendency for repeated statements to receive higher truth ratings because familiarity or processing fluency is mistaken for evidence. Prior knowledge is no automatic shield.

Cognitive dissonance. Discomfort or inconsistency arising among beliefs, self-image and behaviour. People may change an attitude, reinterpret conduct or add justification to restore coherence.

Self-persuasion. Change produced partly by reasons, choices or explanations generated by the recipient. It can create ownership, provided the initial choice is informed and free from concealed escalation.

Go Deeper

Four routes beyond the one-hour model, chosen for different jobs.

The overview

Daniel J. O'Keefe, *Persuasion: Theory and Research*, third edition (2016). The next step for a view of the field rather than another collection of tactics. O'Keefe is unusually careful about definitions, effect sizes, comparison conditions and the difference between what a study found and what a practitioner wishes it found. The book covers source, message, receiver, resistance and behavioural consequences while showing why broad questions such as whether fear works are badly formed. It is an academic text, so the pace is measured and the references are part of the value. It is also the book most likely to make you distrust a universal persuasion rule for the right reason.

The primary source

Aristotle, *The Art of Rhetoric*, translated by Robin Waterfield and edited by Harvey Yunis (2018). Read this to see persuasion before psychology divided it into variables. Aristotle treats rhetoric as judgement under uncertainty and asks what a speaker can reasonably draw from character, emotion and argument in a particular case. The work belongs to a political world unlike ours and includes assumptions that require historical distance. Waterfield's translation is readable, while Yunis's introduction explains the civic and intellectual setting without turning the text into a modern speaking manual. The compressed arguments reward slow reading; this is the least hurried book on the list.

The practical classic

Robert B. Cialdini, *Influence, New and Expanded: The Psychology of Persuasion* (2021). Its practical map covers reciprocity, authority, social proof, liking, scarcity, unity, and the pull of commitment and consistency. Its cases make social influence visible in shops, organisations and everyday requests. Read it with one question in reserve: did the example produce compliance, private attitude change or durable behaviour? The principles are useful mechanisms, but the popular culture around them often removes the conditions and sells buttons. Cialdini is strongest when read as observation, not incantation. The expanded edition adds unity and brings the examples into the digital age.

The counterweight

Hugo Mercier, *Not Born Yesterday: The Science of Who We Trust and What We Believe* (2020). Mercier attacks the flattering fear that other people are gullible vessels waiting for propaganda. Humans possess vigilance mechanisms for judging sources, checking coherence and discounting self-interested claims. The argument can be pushed too far if treated as proof that manipulation is rare or harmless, but it restores a missing half of persuasion: recipients resist for reasons. Read it after Cialdini. One shows recurrent influence; the other asks why influence is difficult and why much attempted manipulation fails. It is a book about the ecology of trust as much as the reception of arguments.

Notes and Sources

Scope and definitions

Persuasion. Definitions vary across communication and social psychology. Daniel O'Keefe's *Persuasion: Theory and Research* is the main guide to the definitional problem, especially the roles of intention, communication and some measure of recipient freedom. This book uses *persuasion* for intentional communication aimed at changing belief, attitude, intention or voluntary behaviour. It separates persuasion from coercion, which changes the consequences of refusal, and from manipulation, which hides material aims, information or pressure. That boundary is an ethical working distinction rather than a claim that every scholar uses the terms identically.

The outcome ladder. Exposure, attention, comprehension, belief, attitude, intention, action and persistence form the book's organising model. It is a synthesis, not a renamed single theory. Hovland, Janis and Kelley provide the early reception and acceptance tradition. Albarracín and Shavitt review the modern attitude literature. Webb and Sheeran quantify the gap between experimental changes in intention and later behaviour. Conner and Norman provide the newest broad synthesis used here. Glasman and Albarracín show why attitude qualities and measurement conditions alter prediction of later behaviour, while Verplanken and Orbell review the role of habit.

The Whole Thing in One Page and Why You Should Care

Conditional effects. The claim that familiar influence principles are conditional rests on the accumulated pattern across persuasion research rather than one nullifying study. Wood's review, O'Keefe's synthesis and Petty and Cacioppo's model all show that source, message and recipient variables can interact. An effect measured on recall, attitude or immediate compliance does not automatically establish a durable effect on behaviour.

Digital and conversational systems. The opening refers to cheap message testing and adaptive dialogue. The current evidence is treated

later and bounded to controlled changes in stated policy attitudes. No claim of mind control, stable superiority or durable mass behaviour change is made.

The Core Ideas

Audience starting conditions. The discussion of attitude strength, accessibility, direct experience and behavioural prediction draws on Glasman and Albarracín, Albarracín and Shavitt, and Wood. The text does not reduce all selective reasoning to identity defence. Accuracy, belonging, self-consistency and practical goals can all shape processing.

Motivational matching. Joyal-Desmarais and colleagues synthesised 702 experimental tests of messages matched to recipient motivations. Their meta-analysis found a positive average matching effect across attitudes, intentions and behaviour, with heterogeneity and possible costs from mismatching. Teeny and colleagues provide the broader framework for personalised matching. The manuscript says only that matching can help and makes variation visible.

Moral reframing. Feinberg and Willer's experiments tested arguments framed in moral values associated with the recipient's political group in United States samples. Their later review describes the approach and its boundary conditions. These studies support the possibility of honest value translation, not a universal recipe or permission to imitate convictions one does not hold.

Familiarity and truth. Fazio and colleagues found that repetition could raise truth ratings even for statements whose correct answers participants knew. Dechêne and colleagues provide a meta-analytic review of the repetition-based truth effect. Montoya and colleagues review the related mere-exposure literature. The manuscript distinguishes familiarity, liking, recall and belief and does not say repetition overrides strong knowledge in every setting.

Credibility. Hovland, Janis and Kelley established source credibility as a central persuasion variable. Pornpitakpan's review found a generally

positive relationship while documenting inconsistency and moderators. Petty and Cacioppo, and later work summarised by O'Keefe, support the claim that a source cue can affect the amount, direction or confidence of thought rather than functioning as a single shortcut. Expertise and trustworthiness remain audience and domain judgements.

Elaboration and cognitive response. Petty and Cacioppo's elaboration likelihood model and Chaiken's heuristic-systematic research are the main sources. The models concern processing conditions and functions, not anatomical routes or rational versus emotional kinds of person. Petty, Cacioppo and Goldman (1981) varied argument quality, attributed expertise and personal relevance among 145 United States undergraduates. The proposed examination would begin the following year or in ten years; high relevance increased the weight of argument quality, while low relevance increased the weight of source expertise. The example illustrates an interaction, not a universal division between evidence and authority.

Narrative persuasion. Braddock and Dillard's meta-analysis supports modest average relationships between narratives and story-consistent beliefs, attitudes, intentions and behaviour. It does not establish that a story generally beats statistics or that narrative effects are always durable. The text therefore recommends joining human consequence to representative evidence.

Fear appeals. Tannenbaum and colleagues synthesised 127 papers, 248 independent samples and 27,372 participants. These included unpublished reports. Fear appeals had a positive average effect on a composite of attitudes, intentions and behaviour, with larger effects when efficacy information was included. The analysis did not identify a general backfire effect when efficacy information was absent. These average findings do not certify every threatening message or a guaranteed behavioural result.

Conformity and norms. Asch's line-judgement studies distinguish public response from private conviction, while showing substantial independence. His 1955 account reports interviews suggesting self-doubt as well as yielding despite misgivings; no single motive explains all conformity. Cialdini

and colleagues distinguish descriptive from injunctive norms. Schultz and colleagues' California household study found that neighbour-comparison electricity feedback could increase consumption among low users, an effect prevented by an added approving cue. The study shows context-sensitive effects, not a universal coefficient. Informative social proof requires scrutiny of independence and selection; persuasive social proof need not be accurate.

Reactance. Dillard and Shen connect perceived threat to freedom with anger and negative cognition in persuasive health messages. Steindl and colleagues review later developments. The manuscript describes reactance as a possible process, not a law that firm language always reverses an effect.

Inoculation. Banas and Rains' meta-analysis supports an average resistance effect from forewarning, weakened challenge and refutational material. Duration, transfer and message design vary. The biological metaphor is used cautiously.

Dissonance and self-persuasion. Festinger's theory and the Festinger-Carlsmith forced-compliance experiment are the historical sources. The 1959 study had a small sample and has generated competing accounts of the mechanism. It is included as a field-shaping anchor rather than a complete proof that behaviour always changes attitude. The broader self-persuasion claim is bounded to reasons, explanations and choices that recipients generate under transparent conditions.

Interpersonal conversation. Broockman and Kalla (2016) and Kalla and Broockman (2020) support durable change on selected exclusionary attitudes after structured conversations; their interventions are treated as packages. Santoro and colleagues tested a listening manipulation in a preregistered longitudinal field experiment. The analysed sample was 1,485 United States participants, not the total number of conversations conducted. Listening added to a fixed persuasive narrative improved relational measures and reduced defensiveness, but produced no detectable extra attitude change immediately or five weeks later. The authors did not test

tailoring the appeal to the information elicited by listening. The manuscript therefore separates that tested result from the practical possibility of using listening diagnostically.

Conversational artificial intelligence. Salvi and colleagues randomly assigned 900 United States participants to short debates with human or GPT-4 opponents, with or without access to six sociodemographic attributes. Personalised GPT-4 significantly outperformed the human-human baseline. An author correction published on 3 September 2026 reported that the direct personalised versus unpersonalised GPT-4 comparison was not statistically significant at the conventional threshold, with $P = 0.07$. The manuscript therefore does not attribute the observed advantage conclusively to personalisation. Hackenburg and colleagues ran three experiments with 76,977 participants, 19 models and 707 political issues, then checked 466,769 model claims. In those experiments, persuasion post-training and prompting increased persuasiveness more than model scale or personalisation, and the same interventions reduced factual accuracy. Chen and colleagues used two survey experiments with 10,417 participants plus simulations. They found chatbot conversations about as persuasive as campaign adverts conditional on exposure. Their simulations favoured traditional methods for mass reach under the modelled reception and engagement assumptions; this was not an observed head-to-head deployment result. These studies concern selected political topics, stated attitudes and defined systems. They do not establish persistent behaviour change, universal superiority or inevitable mass impact.

Excluded AI evidence. Costello, Pennycook and Rand's 2024 *Science* paper on conspiracy-belief dialogue received an Editorial Expression of Concern from H. Holden Thorp in 2026. It is named only to explain its exclusion and supplies no positive claim in the manuscript or bibliography of supporting studies.

Historical and operating sequence

Aristotle. The account of rhetoric, ethos, pathos, logos and the enthymeme draws on *The Art of Rhetoric*, especially the definition of rhetoric as finding the available persuasive means in each case. The text paraphrases rather than quotes. The example of an unstated premise is illustrative of the form and is not presented as a reported Athenian case.

The wartime research programme. Hovland, Lumsdaine and Sheffield reported the United States Army studies of the *Why We Fight* films. The films increased some factual knowledge more reliably than they changed broad attitudes or motivation. The manuscript compresses a complex research programme into that bounded contrast and does not claim all films or outcomes were identical.

The Yale programme. Hovland, Janis and Kelley organised research around communicator, communication, audience predisposition and response. The manuscript's wording about source, sidedness, conclusions, fear and order describes the programme's research questions rather than attributing every experiment to the 1953 volume alone.

Persuasion principles. Cialdini's expanded framework supplies reciprocity, commitment and consistency, social proof, authority, liking, scarcity and unity. The manuscript interprets these as mechanisms or diagnostic questions. The detailed practical examples are illustrations, not claims that each precise scenario was tested in Cialdini's work.

Campaign contact. Kalla and Broockman's 2018 meta-analysis covered 49 field experiments and found an average persuasive effect of campaign contact in United States general elections close to zero. The authors also distinguish primaries and less saturated contexts. The manuscript keeps the population, election type and average explicit and rejects the universal inference that political persuasion is impossible.

Misinformation correction. Ecker and colleagues review misinformation belief, continued influence and correction. Wood and Porter found broad backfire difficult to elicit across many political claims. Walter and Murphy's

meta-analysis supports the general effectiveness of corrective messages while documenting moderators. The practical recommendations to foreground the accurate account, explain the error and offer an alternative model are drawn from this literature but are not represented as a guaranteed sequence.

Adaptive testing. The discussion of digital optimisation is conceptual. It distinguishes optimisation of a measured proxy from the outcome a communicator claims to value. Advertising technology and platform economics remain outside this book's owned scope.

Illustrations. The manager and spreadsheet, service trial, flood defences, train cancellation, colleague and council food-waste examples are hypothetical. The council case is explicitly introduced as such. No household, council, participation rate or treatment effect is reported. It demonstrates the outcome ladder, source trust, norm selection, friction, agency and delayed measurement in one case.

How we know. The limitations paragraph reflects the field's dependence on laboratory attitudes and intentions, the concentration of canonical research in Western populations, and the difficulty of comparing outcomes across domains. Meta-analysis increases precision around an average only to the extent that included measures, populations and designs answer a comparable question.

What People Get Wrong

The 7-38-55 formula. Mehrabian and Wiener's first study concerned inconsistent verbal and vocal communications of attitude. Mehrabian and Ferris studied words, tone and facial cues in judgements of liking. The familiar percentages were derived by combining narrow results. They do not show that words carry 7 per cent of all meaning. The manuscript retains the role of non-verbal information while rejecting the general ratio.

The Vicary story. James Vicary's 1957 cinema claim lacked an adequate documented experiment, failed reproduction and was later admitted to be a promotional fabrication. Pratkanis reviews the history and the wider culture

of subliminal persuasion claims.

Subliminal self-help. Greenwald and colleagues ran double-blind tests of commercial subliminal tapes. Perceived improvement followed the labels attached to tapes rather than their hidden content. Karremans, Stroebe and Claus found a narrow, goal-dependent effect on brand choice among thirsty participants. The text uses the contrast to separate constrained priming from control over complex preferences.

Backfire. The correction myth is addressed with Ecker and colleagues, Wood and Porter, and Walter and Murphy. The manuscript says backfire can occur but is not the normal effect. Continued influence and decaying correction remain visible limitations.

Use It

Practical status. The diagnostic sequence in this section is an application of the book's evidence model, not a validated universal checklist. It asks communicators to define the outcome, identify the blocked stage, make credibility inspectable, invite useful thought, reduce needless social punishment and measure delayed effects. Recommendations involving medical, legal, employment or public decisions still require the relevant professional standards and informed-consent rules.

Measurement. Webb and Sheeran support the warning that changes in intention do not map one-for-one onto behaviour. Glasman and Albarracín support attention to attitude strength and correspondence. The advice to use random assignment where practical follows standard causal inference rather than a persuasion-specific theorem. The required outcome, time horizon and ethical comparison determine the design.

Bibliography

Albarracín, Dolores, and Sharon Shavitt. "Attitudes and Attitude Change." *Annual Review of Psychology* 69 (2018): 299-327. DOI: 10.1146/annurev-psych-122216-011911.

Aristotle. *The Art of Rhetoric*. Translated by Robin Waterfield. Edited by Harvey Yunis. Oxford: Oxford University Press, 2018.

Asch, Solomon E. "Opinions and Social Pressure." *Scientific American* 193, no. 5 (1955): 31-35. DOI: 10.1038/scientificamerican1155-31.

Asch, Solomon E. "Studies of Independence and Conformity: I. A Minority of One Against a Unanimous Majority." *Psychological Monographs: General and Applied* 70, no. 9 (1956): 1-70. DOI: 10.1037/h0093718.

Banas, John A., and Stephen A. Rains. "A Meta-Analysis of Research on Inoculation Theory." *Communication Monographs* 77, no. 3 (2010): 281-311. DOI: 10.1080/03637751003758193.

Braddock, Kurt, and James Price Dillard. "Meta-Analytic Evidence for the Persuasive Effect of Narratives on Beliefs, Attitudes, Intentions, and Behaviors." *Communication Monographs* 83, no. 4 (2016): 446-467. DOI: 10.1080/03637751.2015.1128555.

Broockman, David E., and Joshua L. Kalla. "Durably Reducing Transphobia: A Field Experiment on Door-to-Door Canvassing." *Science* 352, no. 6282 (2016): 220-224. DOI: 10.1126/science.aad9713.

Chaiken, Shelly. "Heuristic Versus Systematic Information Processing and the Use of Source Versus Message Cues in Persuasion." *Journal of Personality and Social Psychology* 39, no. 5 (1980): 752-766. DOI: 10.1037/0022-3514.39.5.752.

Chen, Zhongren, Joshua Kalla, Quan Le, Shinpei Nakamura-Sakai, Jasjeet Sekhon, and Ruixiao Wang. "A Framework to Assess the Persuasion Risks Large Language Model Chatbots Pose to Democratic Societies." *Journal of Experimental Political Science* (2026): 1-14. Published online 4 June 2026. DOI: 10.1017/XPS.2026.10032.

Cialdini, Robert B. *Influence, New and Expanded: The Psychology of Persuasion*. New York: Harper Business, 2021.

Cialdini, Robert B., Linda J. Demaine, Brad J. Sagarin, Daniel W. Barrett,

- Kelton Rhoads, and Patricia L. Winter. "Managing Social Norms for Persuasive Impact." *Social Influence* 1, no. 1 (2006): 3-15. DOI: 10.1080/15534510500181459.
- Conner, Mark, and Paul Norman. "Attitudes, Intentions, and Behavior Change." *Annual Review of Psychology* 77 (2026): 311-337. DOI: 10.1146/annurev-psych-013125-042110.
- Dechêne, Alice, Christoph Stahl, Jochim Hansen, and Michaela Wänke. "The Truth About the Truth: A Meta-Analytic Review of the Truth Effect." *Personality and Social Psychology Review* 14, no. 2 (2010): 238-257. DOI: 10.1177/1088868309352251.
- Dillard, James Price, and Lijiang Shen. "On the Nature of Reactance and Its Role in Persuasive Health Communication." *Communication Monographs* 72, no. 2 (2005): 144-168. DOI: 10.1080/03637750500111815.
- Ecker, Ullrich K. H., Stephan Lewandowsky, John Cook, Philipp Schmid, Lisa K. Fazio, Nadia Brashier, Panayiota Kendeou, Emily K. Vraga, and Michelle A. Amazeen. "The Psychological Drivers of Misinformation Belief and Its Resistance to Correction." *Nature Reviews Psychology* 1 (2022): 13-29. DOI: 10.1038/s44159-021-00006-y.
- Fazio, Lisa K., Nadia M. Brashier, B. Keith Payne, and Elizabeth J. Marsh. "Knowledge Does Not Protect Against Illusory Truth." *Journal of Experimental Psychology: General* 144, no. 5 (2015): 993-1002. DOI: 10.1037/xge0000098.
- Feinberg, Matthew, and Robb Willer. "From Gulf to Bridge: When Do Moral Arguments Facilitate Political Influence?" *Personality and Social Psychology Bulletin* 41, no. 12 (2015): 1665-1681. DOI: 10.1177/0146167215607842.
- Feinberg, Matthew, and Robb Willer. "Moral Reframing: A Technique for Effective and Persuasive Communication Across Political Divides." *Social and Personality Psychology Compass* 13, no. 12 (2019): e12501. DOI: 10.1111/spc3.12501.
- Festinger, Leon. *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford

University Press, 1957.

Festinger, Leon, and James M. Carlsmith. "Cognitive Consequences of Forced Compliance." *Journal of Abnormal and Social Psychology* 58, no. 2 (1959): 203-210. DOI: 10.1037/h0041593.

Glasman, Laura R., and Dolores Albarracín. "Forming Attitudes That Predict Future Behavior: A Meta-Analysis of the Attitude-Behavior Relation." *Psychological Bulletin* 132, no. 5 (2006): 778-822. DOI: 10.1037/0033-2909.132.5.778.

Greenwald, Anthony G., Eric R. Spangenberg, Anthony R. Pratkanis, and Jay Eskenazi. "Double-Blind Tests of Subliminal Self-Help Audiotapes." *Psychological Science* 2, no. 2 (1991): 119-122. DOI: 10.1111/j.1467-9280.1991.tb00112.x.

Hackenburg, Kobi, Ben M. Tappin, Luke Hewitt, Ed Saunders, Sid Black, Hause Lin, Catherine Fist, Helen Margetts, David G. Rand, and Christopher Summerfield. "The Levers of Political Persuasion with Conversational Artificial Intelligence." *Science* 390, no. 6777 (2025): eaea3884. DOI: 10.1126/science.aea3884.

Hovland, Carl I., Irving L. Janis, and Harold H. Kelley. *Communication and Persuasion: Psychological Studies of Opinion Change*. New Haven, CT: Yale University Press, 1953.

Hovland, Carl I., Arthur A. Lumsdaine, and Fred D. Sheffield. *Experiments on Mass Communication*. Princeton, NJ: Princeton University Press, 1949.

Joyal-Desmarais, Keven, Alexandra K. Scharmer, Molly K. Madzelan, Jolene V. See, Alexander J. Rothman, and Mark Snyder. "Appealing to Motivation to Change Attitudes, Intentions, and Behavior: A Systematic Review and Meta-Analysis of 702 Experimental Tests of the Effects of Motivational Message Matching on Persuasion." *Psychological Bulletin* 148, no. 7-8 (2022): 465-517. DOI: 10.1037/bul0000377.

Kalla, Joshua L., and David E. Broockman. "The Minimal Persuasive Effects of Campaign Contact in General Elections: Evidence from 49 Field

Experiments." *American Political Science Review* 112, no. 1 (2018): 148-166. DOI: 10.1017/S0003055417000363.

Kalla, Joshua L., and David E. Broockman. "Reducing Exclusionary Attitudes through Interpersonal Conversation: Evidence from Three Field Experiments." *American Political Science Review* 114, no. 2 (2020): 410-425. DOI: 10.1017/S0003055419000923.

Karremans, Johan C., Wolfgang Stroebe, and Jasper Claus. "Beyond Vicary's Fantasies: The Impact of Subliminal Priming and Brand Choice." *Journal of Experimental Social Psychology* 42, no. 6 (2006): 792-798. DOI: 10.1016/j.jesp.2005.12.002.

Mehrabian, Albert, and Susan R. Ferris. "Inference of Attitudes from Nonverbal Communication in Two Channels." *Journal of Consulting Psychology* 31, no. 3 (1967): 248-252. DOI: 10.1037/h0024648.

Mehrabian, Albert, and Morton Wiener. "Decoding of Inconsistent Communications." *Journal of Personality and Social Psychology* 6, no. 1 (1967): 109-114. DOI: 10.1037/h0024532.

Mercier, Hugo. *Not Born Yesterday: The Science of Who We Trust and What We Believe*. Princeton, NJ: Princeton University Press, 2020.

Montoya, R. Matthew, Robert S. Horton, Jacqueline L. Vevea, Martyna Citkowicz, and Elissa A. Lauber. "A Re-Examination of the Mere Exposure Effect: The Influence of Repeated Exposure on Recognition, Familiarity, and Liking." *Psychological Bulletin* 143, no. 5 (2017): 459-498. DOI: 10.1037/bul0000085.

O'Keefe, Daniel J. *Persuasion: Theory and Research*. 3rd ed. Thousand Oaks, CA: Sage, 2016.

Petty, Richard E., and John T. Cacioppo. *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*. New York: Springer-Verlag, 1986.

Petty, Richard E., John T. Cacioppo, and Rachel Goldman. "Personal

Involvement as a Determinant of Argument-Based Persuasion.” *Journal of Personality and Social Psychology* 41, no. 5 (1981): 847-855. DOI: 10.1037/0022-3514.41.5.847.

Pornpitakpan, Chanthika. “The Persuasiveness of Source Credibility: A Critical Review of Five Decades' Evidence.” *Journal of Applied Social Psychology* 34, no. 2 (2004): 243-281. DOI: 10.1111/j.1559-1816.2004.tb02547.x.

Pratkanis, Anthony R. “The Cargo-Cult Science of Subliminal Persuasion.” *Skeptical Inquirer* 16, no. 3 (1992): 260-272.

Salvi, Francesco, Manoel Horta Ribeiro, Riccardo Gallotti, and Robert West. “On the Conversational Persuasiveness of GPT-4.” *Nature Human Behaviour* 9 (2025): 1645-1653. DOI: 10.1038/s41562-025-02194-6.

Salvi, Francesco, Manoel Horta Ribeiro, Riccardo Gallotti, and Robert West. “Author Correction: On the Conversational Persuasiveness of GPT-4.” *Nature Human Behaviour* (2026). Published online 3 September 2026. DOI: 10.1038/s41562-026-02588-0.

Santoro, Erik, David E. Broockman, Joshua L. Kalla, and Roni Porat. “Listen for a Change? A Longitudinal Field Experiment on Listening's Potential to Enhance Persuasion.” *Proceedings of the National Academy of Sciences* 122, no. 8 (2025): e2421982122. DOI: 10.1073/pnas.2421982122.

Schultz, P. Wesley, Jessica M. Nolan, Robert B. Cialdini, Noah J. Goldstein, and Vladas Griskevicius. “The Constructive, Destructive, and Reconstructive Power of Social Norms.” *Psychological Science* 18, no. 5 (2007): 429-434. DOI: 10.1111/j.1467-9280.2007.01917.x.

Steindl, Christina, Eva Jonas, Sandra Sittenthaler, Eva Traut-Mattausch, and Jeff Greenberg. “Understanding Psychological Reactance: New Developments and Findings.” *Zeitschrift für Psychologie* 223, no. 4 (2015): 205-214. DOI: 10.1027/2151-2604/a000222.

Tannenbaum, Melanie B., Justin Hepler, Rick S. Zimmerman, Lindsey Saul, Samantha Jacobs, Kristina Wilson, and Dolores Albarracín. “Appealing to

Fear: A Meta-Analysis of Fear Appeal Effectiveness and Theories.” *Psychological Bulletin* 141, no. 6 (2015): 1178-1204. DOI: 10.1037/a0039729.

Teeny, Jacob D., Joseph J. Siev, Pablo Briñol, and Richard E. Petty. “A Review and Conceptual Framework for Understanding Personalized Matching Effects in Persuasion.” *Journal of Consumer Psychology* 31, no. 2 (2021): 382-414. DOI: 10.1002/jcpy.1198.

Thorp, H. Holden. “Editorial Expression of Concern.” *Science* 392, no. 6803 (2026): 1131. DOI: 10.1126/science.aej2383.

Verplanken, Bas, and Sheina Orbell. “Attitudes, Habits, and Behavior Change.” *Annual Review of Psychology* 73 (2022): 327-352. DOI: 10.1146/annurev-psych-020821-011744.

Walter, Nathan, and Sheila T. Murphy. “How to Unring the Bell: A Meta-Analytic Approach to Correction of Misinformation.” *Communication Monographs* 85, no. 3 (2018): 423-441. DOI: 10.1080/03637751.2018.1467564.

Webb, Thomas L., and Paschal Sheeran. “Does Changing Behavioral Intentions Engender Behavior Change? A Meta-Analysis of the Experimental Evidence.” *Psychological Bulletin* 132, no. 2 (2006): 249-268. DOI: 10.1037/0033-2909.132.2.249.

Wood, Thomas J., and Ethan Porter. “The Elusive Backfire Effect: Mass Attitudes' Steadfast Factual Adherence.” *Political Behavior* 41 (2019): 135-163. DOI: 10.1007/s11109-018-9443-y.

Wood, Wendy. “Attitude Change: Persuasion and Social Influence.” *Annual Review of Psychology* 51 (2000): 539-570. DOI: 10.1146/annurev.psych.51.1.539.