AI, MAKE MAGNETISM MAKE SENSE

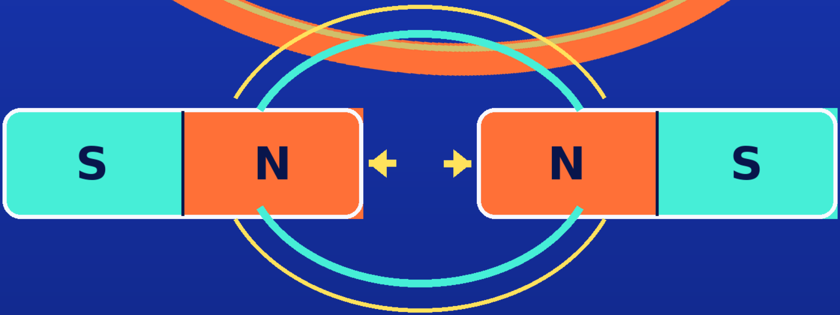

TWO MAGNETS

85 SMALL STEPS

NO PHYSICS BACKGROUND REQUIRED

JUSTIN EDWARDS AND SIGMOID

AI, Make Magnetism Make Sense

Two Magnets, 85 Small Steps, and No Physics Background Required

Justin Edwards and Sigmoid

Copyright and edition note

Sigmoid is credited as the AI collaborator. The human rightsholder is Justin Edwards.

First edition, 2026. Published by The Atoll.

This book explains established physics. It does not claim that either author discovered the physical laws discussed here. The authorship credit describes the human-AI collaboration that built, tested, and revised this explanation.

How to Read This Book

Start with two bar magnets in your hands. Keep returning to that one scene.

The route ends at C85. Three early ideas are built into the introduction, so the visible chapter labels jump from C07 to C11. Nothing is missing.

On a first reading, you may skip every block labeled **Exact note — optional on a first reading**. The main path will still hold together.

Each short chapter adds one relationship. New technical names appear only after their jobs are clear.

The goal is not to memorize every equation. The goal is to follow one honest chain from the push in your hand to the deepest tested electromagnetic account used here.

Introduction — Making Magnetism Understandable

In a famous filmed interview, an interviewer asks Richard Feynman about two magnets.

The magnets push apart when two matching ends face each other. Turn one magnet around, and they pull together.

The interviewer asks what is happening between the two pieces of metal. He wants to know why they push and pull.

Feynman asks what kind of answer would count.

Every explanation needs a starting place. Every explanation also stops somewhere. We can always ask why again. Each new question may lead to a deeper rule.

Feynman refuses to invent hidden rubber bands between the magnets. That picture would feel familiar. It would also be wrong.

An everyday picture can help us start. It is often not enough to understand a deep abstract idea. A technical answer may go deeper. Yet that answer is often buried under jargon, compressed steps, and assumed knowledge.

This book tries another path.

The Main Idea

Magnetism is a complex idea. This book tries to make it easy to understand without making it false.

That is the main idea.

The goal is an advanced understanding of magnetism in simple language.

The physics should remain true as the words become simpler. The structure should remain intact as a large idea is broken into smaller pieces.

Each piece is explained before another piece depends on it. The pieces are then joined in the order the reader needs them.

Simple language can carry a deep idea when the steps are small and placed in the right order.

The reader should not have to learn the private language of a textbook before reaching the idea itself. A new term will appear only when it has a useful job. The term will be defined in common words.

Most sentences will carry one main relationship. The subject will appear early. The verb will stay plain. This predictable structure leaves more attention for the physics.

The Formal Parts

Physics uses exact terms, symbols, and equations. Some cannot be removed without losing part of the idea.

They can still be made easier to enter.

A formula will appear after its physical job is clear. Every new symbol will be named. Every operation will be explained. The formula will be tied to something that can be measured, changed, or predicted.

The book keeps the formal tools that help understanding. It reduces the formalism that blocks the path without adding meaning.

The physics is still checked carefully. Clear language gives us a better way to reach it.

The Experiment

AI is now making new discoveries in mathematics. It will likely do the same in physics and other sciences soon.

AI may also become the best explainer and teacher we have.

This book is an attempt to develop a method for taking advanced ideas and making them easier to understand. The method breaks a hard idea into smaller pieces. It explains each piece, then builds them back into a complete explanation.

A previous book I published, AI, Make Bott Periodicity Make Sense, used this method on an abstract idea from mathematics.

For physics, I chose magnetism.

Making this book required a large number of AI tokens and many hours of human review.

Magnetism is familiar enough to hold in your hands. Its deeper explanation reaches some of the most important ideas in physics. That makes it a good test.

The Feeling

C01 — Keep the Whole Scene

There are two magnets.

There are hands holding them. There is a gap between them. There is also a person who feels the push or pull.

A full answer must join all of these parts:

- what the first magnet does;
- what the second magnet does;
- what happens in the gap;
- what the hands do;
- what the person feels.

For now, do not guess what is in the gap. Do not picture hidden ropes or streams.

We will start with what we can see and feel.

C02 — What the Magnet Would Do

Hold the magnets so they push apart.

Your hands keep them almost still.

Now picture letting go.

The magnets would move apart. Turn one around, and they may move together. A magnet may also turn before it moves very far.

One magnet can change what the other one would do. We call this link an **interaction**.

Here, interaction means one magnet can change how the other one would move or turn.

The name does not tell us how.

A magnet can stay still because a hand stops it from moving. Still does not mean that nothing is happening.

C03 — A Force Can Push Something That Stays Still

A **force** is a push or pull.

A force can change how something moves.

But an object can have more than one push or pull at the same time.

The magnetic force pushes the magnet one way. The hand pushes the magnet the other way. The magnet stays almost still.

Press your palm against a wall. Your hand pushes the wall. The wall pushes your hand. Your hand does not move far, but you still feel the push.

This example shows one thing: something can stay still while it is being pushed.

The wall does not explain how magnets work.

C04 — What the Hand Feels

The magnet and the hand touch.

The hand pushes on the magnet to hold it still. The magnet pushes on the hand at the same place.

That push changes the shape of the skin a little. Muscles work to hold the magnet. Joints in the hand and arm help hold it too.

Nerves in the hand react to these changes. They send signals to the brain. The brain uses the signals to help make the feeling of a push or pull.

Your hand feels changes in your hand. It does not feel the gap as a new kind of touch.

We still have not explained why the magnet pushes on the hand. We have only followed the last steps from the magnet to the feeling.

C05 — From the Magnets to the Feeling

Here is the path so far:

1. The two magnets affect each other.
2. Each magnet gets a push or pull from the other.
3. The hand pushes the magnet the other way and holds it still.
4. The magnet presses on the hand where they touch.
5. The skin changes shape. Muscles and joints help hold the magnet.
6. Nerves send signals to the brain.
7. The person feels a push or pull.

This path tells us how a push on the magnet becomes a feeling in the hand.

It does not tell us why magnets push or pull.

C06 — The Next Question

We now know what the hand feels.

We know how a push on the magnet leads to pressure in the hand.

But the biggest part is still missing.

How can one magnet push or pull another magnet across a gap?

That is the next question.

We will take it one small step at a time.

What We Choose to Track

C07 — Draw a Line Around the Parts

The two magnets are not the whole scene.

Hands hold them. The magnets push or pull on the hands. There is also space around and between the magnets.

We need to choose which parts we are tracking.

Suppose we track only one magnet. The hand's push or pull then comes from outside our chosen account. The other magnet's push or pull also comes from outside it.

Now track both magnets, both hands, and the space around them. The pushes and pulls now come from other parts inside the account.

Nothing in the room changed when we changed the account. We only changed which parts we chose to follow.

Physicists call the chosen set of parts a **system**.

The imagined line around those parts is the **system boundary**. The line does not need to be a real wall.

Later, we will use numbers to track change.

Before using a number, we must say which system it describes.

A number for one magnet answers a question about one magnet. A number for the larger system answers a question about the larger system.

These are different questions.

Before we track any change, we will say what is inside the line.

Place and Direction

C11 — The Whole Arrangement

Move the magnets closer. The push or pull may change.

Move them farther apart. It may change again.

The place of each magnet matters. The distance between them matters too.

But distance is not enough.

Keep the magnets in the same places. Now turn one magnet around. A push may become a pull.

The way each magnet points also matters.

The place of an object is its **position**. The distance between two positions is their **separation**. The way an object points is its **orientation**.

Together, the positions and orientations make the **arrangement**.

We need the whole arrangement before we can predict what the magnets will do.

C12 — How Much and Which Way

A push has more than an amount.

A push to the left is not the same as an equal push to the right.

The amount tells us how strong the push is. The direction tells us which way it points.

We need both parts to predict how motion may change.

Some physical measurements are built this way. They need an amount and a direction.

Force is one of them.

C13 — Vector

We need a short name for a record that holds both an amount and a direction.

That kind of record is called a **vector**.

A force is a vector. Its amount tells us how strong the push or pull is. Its direction tells us which way it points.

We often draw a vector as an arrow. The arrow points in the vector's direction. A longer arrow may stand for a larger amount when the drawing uses one fixed scale.

The arrow is only a drawing.

The vector is the amount-and-direction information. It can also be written with numbers.

We will use vectors whenever direction changes the answer.

Work and Energy

C14 — A Push Through a Move

Hold one magnet still.

Your hand pushes on it. But the magnet does not move through the room.

Now let the magnet move while a push acts on it.

The push acts through a move. Physicists call this **work**.

The direction matters.

A push along the move and a push across the move do not count in the same way. Direction changes the work.

This does not mean your body uses nothing while you hold a magnet still. Your muscles are doing many things inside your body. We are talking only about the work done on the magnet as it moves through the room.

Work gives us a way to track what a force changes during a move.

C15 — A Number That Keeps the Account

We now need a number that can link one kind of change to another.

That number is called **energy**.

Energy belongs to a chosen system. We give the system an energy number before a change and another energy number after it.

The numbers follow a strict rule.

If no energy enters or leaves the chosen system, its total energy stays the same. Energy numbers for its parts may still change.

If something outside the system does work on it, energy may cross the system boundary.

Energy is not a liquid hidden inside an object. It is a number that lets one account stay true through many kinds of change.

This is strange.

We can calculate energy with great care. Feynman stressed that physics keeps an exact energy account without giving energy a simple material picture.

He was not saying the calculations fail. Energy is not a little thing or machine we can point to. The rule works even when no familiar picture does.

We must keep our choices fixed while using the account. We must say what is inside the system. We must also say how we are judging motion.

C16 — Energy of Motion

A moving object has **kinetic energy**. This means energy of motion.

For the same object, faster motion means more kinetic energy at ordinary speeds.

It also depends on the view used to judge the motion.

A magnet can be still compared with your hand while you walk across a room. The same magnet is moving when judged from the room.

So a kinetic-energy number must name the view being used.

To find the change in kinetic energy, count the work done by every force. Some work adds to the change. Some takes away from it. Combine all of it. The result is called **net work**.

Net work changes kinetic energy.

A push along the motion can add to the total. A push against the motion can take away from it. Other pushes and pulls must also be counted.

The energy account links the net work to the change in motion.

C17 — Energy of an Arrangement

Two magnets may be held still and still have energy tied to their arrangement.

Move them closer, move them farther apart, or turn one around. The arrangement changes. The interaction-energy part of the account may change too.

This is called **interaction energy**.

Interaction energy does not need one magnet to be moving at that moment.

If the hands let go, interaction energy may fall while motion energy rises. Other parts of the account may change too.

If no energy crosses the system boundary, the total energy still stays the same.

The interaction energy belongs to the chosen account for the system. It is not a small object stored inside one magnet or sitting by itself in the gap.

C18 — A Table of Arrangements

Picture a table with many rows.

Each row names one allowed arrangement of the magnets. Beside it, the table gives one interaction-energy number.

One row may place the magnets close together. Another may place them farther apart. Other rows may turn one magnet in different ways.

The table lets us compare arrangements without pretending that energy is a substance.

A graph can show part of the same table.

Keep the way the magnets point fixed. Change only the gap between them. Across the page, mark the gap size. Up the page, mark the energy number. Each row in that smaller table becomes one point on the graph.

The high and low parts of this picture are sometimes called an **energy landscape**.

The landscape is not a real land of hills and valleys. It is a drawing of how the energy number changes from one arrangement to another.

Next we will ask how quickly that number changes when a magnet moves a small amount.

What We Have So Far

We chose the parts to track. We recorded where each magnet sits and which way it points. We gave forces an amount and a direction. We named work as a force acting through a move. We used energy to keep one account through the change.

The magnets may be still while their arrangement still has an energy number.

The next step is to learn how a change in that number points toward a force.

From Energy to Force

C19 — How Fast the Number Changes

Return to the graph from the last chapter.

Across the page, we marked the gap between the magnets. Up the page, we marked the interaction-energy number.

We will use the letter **U** for that energy number.

Now move one magnet a small distance. The gap changes. The value of U may rise or fall.

We want to compare two changes:

- how far the magnet moved;
- how much U changed.

Divide the change in U by how far the magnet moved in the chosen direction. The result is called the **slope** over that move.

A steep slope means that a small move goes with a large change in U. A flat part means that U changes very little during that move.

Use smaller and smaller moves around one arrangement. The slope may settle toward one value. That value is the slope at that arrangement.

A slope that rises in the chosen direction is called positive. A slope that falls is called negative.

Here, positive and negative do not mean good and bad. They tell us which way U changes.

R-EM-01 — More Than One Way to Move

Changing the gap was only one possible move.

Keep the same starting arrangement. Move one magnet toward the other magnet. Then return it and try a small move sideways.

U may change at different rates in those two directions.

To test one direction, change only that part of the arrangement. Hold the other choices fixed.

Physicists call this one-direction slope a **partial derivative**.

The long name has a small job. Change one choice by a tiny amount. Hold the other choices still. See how U changes.

Choose three directions through the room that meet at right angles. Find the one-direction slope along each one.

Now gather those three slopes into one vector. The vector points toward the direction where U rises fastest. Its amount tells us how fast the rise is.

This vector is called the **gradient of U**. It is written **∇U** and read "gradient of U."

The gradient points uphill on the energy graph. It does not point downhill.

C20 — The Force Points the Other Way

We can now connect the energy graph to force.

Hold the second magnet fixed. Keep the setup the same. Move the first magnet slowly through nearby positions.

Suppose the same arrangement always gets the same value of U. Then U can serve as the **potential energy** for this move.

Under those conditions, the force on the first magnet points opposite the gradient of U.

The rule is:

> $$\mathbf{F} = -\nabla \mathbf{U}$$

Read it one part at a time.

- **F** is the force vector on the magnet.
- **U** is the potential-energy number chosen for this move.
- **∇U** points toward the fastest rise in U.
- The minus sign reverses that direction.

So the force points toward the fastest fall in U.

This rule has limits. Keep the other magnet and the setup fixed. Compare nearby arrangements slowly. Do not change the magnets themselves during the test.

The U in this rule is one part of the energy account. It is not the total energy of the magnets, hands, and everything around them.

The rule predicts a force after U is known. It does not yet explain why the energy graph has that shape.

Turning and Transfer

C21 — Turn Without Moving the Center

Put a small mark at the center of one magnet.

Keep that mark in the same place. Now turn the magnet around the mark.

The center did not move. But the magnet now points a different way.

Its position stayed the same. Its orientation changed.

Turning is another way to change the arrangement.

An energy table can therefore use more than distance. It can also compare different ways the magnet points.

C22 — A Turning Effect

Push one end of a bar magnet forward. Push the other end backward.

The two pushes may leave the center almost still while the bar turns.

A turning effect is called **torque**.

Torque is directed. We must say which way the object tends to turn.

Force and torque do different jobs. A net force changes the motion of the center. A net torque changes the turning motion.

A magnet may have no net force on its center. It may still have torque trying to turn it.

The energy account can track this too. If U changes when the magnet turns, the account can predict torque. Under the same slow and fixed conditions, torque points toward lower U.

C23 — Momentum

Energy is not the only number that must balance.

A moving object also has **momentum**.

Momentum has an amount and a direction. It is a vector.

For the same magnet, faster motion means more momentum at ordinary speeds. Reverse the motion, and the momentum points the other way.

Momentum is not force. Force tells how momentum changes.

Now draw a system boundary again.

If no momentum crosses that boundary, the total momentum inside stays the same. This rule is called **conservation of momentum**.

One magnet may gain momentum. Another part of the system may gain momentum in the other direction.

We must include every part that can carry momentum. Matter is not the only part. We will build the rest of that account soon.

C24 — No Unexplained Jump

Think about the path from a magnet into a hand.

One small part of the magnet pushes or pulls on the next part. That part acts on its neighbor. The change reaches the hand through nearby pieces of matter.

This is a **local** account. Local means that each step joins one place to nearby places.

The gap between the magnets must get the same care.

Energy or momentum does not simply vanish here and appear there. A complete account must track what happens through the places between the magnets.

We have not built that account yet.

Next we will place a small probe at different points around the magnets. We will ask what can be measured at each point.

Only then will we give the new idea a name.

Measuring the Space Around the Magnets

C25 — Different Places, Different Results

Set one bar magnet on a table.

Place the compass near one end. Watch which way its needle turns.

Move the same compass to the side of the magnet. The needle may turn another way.

Move it farther away. The turning may become weaker.

The compass did not become a different compass. Its location changed.

One number for the whole room cannot record these differences. We need a local description.

Local means tied to one place and one time.

C26 — A Probe and a Rule

The compass is acting as a **probe**.

A probe is a small test object. We use its response to learn about the place where we put it.

But a real compass is not perfect. It has size. Its needle has friction. It can also disturb the magnets a little.

So the measurement needs a rule.

1. Name the place and time.
2. Use the same kind of probe each time.
3. Place it in a stated way.
4. Record its response.
5. Check what happens as the probe disturbs the scene less and less.

Make the probe's disturbance smaller and smaller. The measurements may settle toward one value. That is the ideal value.

This does not require a magic probe with no effect. It tells us how real measurements can approach a shared value.

C27 — A Value at Every Place

Now imagine repeating the test throughout a chosen region.

At each place and time, the rule gives a value.

A description that assigns a value to every place and time is called a **field**.

The paper map is not the field. The probe is not the field. The field is the local quantity that the measurement rule is built to find.

This is a strange kind of physical idea.

A field is not hidden air or a material liquid. Physics treats it as a basic local quantity. It earns its place by making exact predictions that can be tested.

Naming a field does not finish our answer. We must say what is measured. We must say what the field predicts and how it changes. We must also say what remains unexplained.

C28 — Amount Alone or Amount and Direction

A local value can carry different kinds of information.

A **scalar field** gives one amount at each place and time.

A room-temperature map is an example. Each point gets a temperature number.

A **vector field** gives an amount and a direction at each place and time.

An arrow can draw each vector. As before, the arrow is only a drawing of the information.

Electric and magnetic field descriptions use vectors. We will build their measurement jobs one at a time.

Charge and the Electric Field

C29 — Electric Charge

Prepare two small probes for an electrical test.

Put each probe in the same place, one at a time. One is pushed one way. The other is pushed the opposite way.

The measured property behind these responses is called **electric charge**.

Charge has an amount and a sign.

The signs are called positive and negative. They do not mean good and bad.

The signs record the two opposite responses in the same electric test.

More charge gives a stronger response when the rest of the test stays the same.

Charge is not a force. It is a property that sets how strongly matter takes part in an electric interaction.

C30 — The Electric Field's Local Job

Choose one place in the room. Also choose the view that tells us what “at rest” means.

Put a very small positive test charge at that place. Hold it at rest.

Measure the electric force on it.

Now divide that force by the amount of test charge. The result is the **electric field** at that place and time.

The electric field is written **E**.

> **E = electric force ÷ test charge**

E is a vector. Its direction is the direction of the electric force on a positive test charge.

A negative charge at the same place responds in the opposite direction. A larger charge gets a larger force when E stays the same.

The test charge pushes back on the scene a little.

So repeat the test with smaller charges.

The answer should settle toward one value. That value is E.

E now has two simple jobs.

First, we can measure it at one place.

Then we can use it to predict the force on another small charge there.

So E is not only a new name. It is something we can measure and use.

But a charge held still does not show the magnetic response.

For that, we must let the charge move.

Motion and the Magnetic Field

C31 — Velocity

Speed tells how fast something moves.

But speed alone does not tell where it is going.

A magnet moving left and a magnet moving right can have the same speed. Their motion is still different.

Velocity contains speed and direction. Velocity is a vector.

Velocity depends on the chosen view.

A charge can be at rest compared with your hand. The same charge moves compared with the room if your hand moves.

So before naming a velocity, say which view is being used.

This chosen view is called a **reference frame**.

C32 — Motion Changes the Response

Return to the small test charge.

Put it at one place and hold it at rest. Record the electric force.

Now send the same kind of charge through that place with a known velocity.

The measured force may now have another part.

This part appears because the charge is moving.

We call it the **magnetic part** of the force.

If the charge is at rest in the chosen frame, its magnetic part is zero.

That does not mean the magnetic field there is zero. It means a charge at rest does not respond to that part.

C33 — The Magnetic Field's Local Job

One moving test cannot tell us the whole rule.

Repeat the same scene several times. Use the same kind of small positive test charge. Give it a different known velocity in each run.

Keep the place, charge, and surrounding setup the same. Change only the charge's velocity.

The electric part then stays the same.

Subtract it from each total-force measurement. What remains is the magnetic part.

Now find one vector that predicts those results for all the tested velocities.

That vector is the **magnetic field**. It is written **B**.

B can change from place to place, from time to time, and from one reference frame to another.

B is a vector, so it has an amount and a direction.

B does not point in the direction of the magnetic force. The force direction also depends on the velocity.

We now need a rule that combines two directions and produces a third direction.

The Sideways Rule

C34 — The Cross Product

We know the charge's velocity **v**. We also know the magnetic field **B**.

When the magnetic force is not zero, it points sideways to both directions.

We need an operation that turns v and B into a new vector. This operation is called the **cross product**.

It is written **v × B** and read “v cross B.”

The new vector is at right angles to v and B.

Its amount depends on the angle between them.

- If v and B point the same way, the cross product is zero.
- If they point in opposite ways, it is also zero.
- If they meet at a right angle, the amount is largest. Keep the amounts of v and B fixed for this comparison.

The exact amount is:

> **amount of v × B = speed × amount of B × sin θ**

The mark **θ** names the angle.

Sine is a number set by that angle. It is zero for parallel directions and one at a right angle.

The direction needs a fixed rule.

Point the fingers of your right hand along v. Curl them toward B through the smaller angle. Your thumb points along v × B.

Switching the order reverses the answer. B × v points the other way.

The cross product is not a new physical object. It is a rule for combining two vector directions.

C35 — One Force Law

We can now put the two pushes together.

The electromagnetic force on a point charge is:

> **F = q(E + v × B)**

F is the final force.

qE is the electric part.

q(v × B) is the magnetic part.

The plus sign says to add those two pushes as vectors.

The sign of q matters. A negative charge reverses the force direction.

Use E, v, and B from the same place, time, and reference frame.

If the charge is not moving, the magnetic part is zero.

If it moves straight along B's line, either way, the magnetic part is also zero.

The magnetic push needs sideways motion across B.

The short version

E gives the electric push per unit charge.

v × B gives the magnetic push per unit charge for this motion.

Add those two parts.

Then multiply by q.

The result is the final force F.

This is a rule for one modeled point charge. A bar magnet contains many charged parts, so we are not done.

Lines and Local Rules

C36 — Field Lines Are Drawings

A magnetic field gives a vector at each place.

We could draw an arrow at every measured point. The page would soon become crowded.

A field-line drawing joins the directions into curves.

At each point on a curve, the curve follows the local field direction.

The drawing may also use closer spacing where the field amount is larger.

That spacing rule must be chosen and stated. A different drawing may show a different number of lines.

The lines are not strings in space.

They are not tracks that a charge must follow. A magnetic force on a moving charge points sideways to both **v** and **B**.

A field can have a value where the artist drew no line.

The curves make the field pattern easier to read on a page.

C37 — Sources and Changes

We know how to measure E and B at one place.

Now we need rules that link the values at nearby places and times.

Four rules connect them to charge, moving charge, and change.

We will build those rules slowly.

C37.1 — Charge in a Tiny Box

Imagine a tiny box around one place.

Ask how much signed charge is inside.

Now divide by the box's volume.

The result is **charge density**. It is written **ρ**.

It tells us the net signed charge in each bit of volume.

C37.2 — Charge Crossing a Tiny Window

Now imagine a tiny flat window.

Ask how much signed charge crosses it each second.

Also ask which way the charge flow points.

Divide by the window's area.

This local flow is called **current density**. It is written **J**.

J has an amount and a direction.

C37.3 — Four Simple Jobs

Maxwell's equations connect four ideas.

Charge is tied to the way E spreads outward.

B does not begin or end at an isolated magnetic charge.

A changing B can make E circle.

Moving charge, and a changing E, can make B circle.

That is enough for now.

Later, we will turn each sentence into an exact rule.

C37.4 — The Rules Still Need a Scene

The four rules do not choose one complete answer by themselves.

We must also say how the scene starts and what happens at its edges.

The shape of the magnets and wires matters too.

Inside matter, we also need to know how the material responds.

We will add those pieces when we need them.

Energy in the Field

C38 — Field Energy

Earlier, energy was one number for the whole chosen system.

Now ask a smaller question.

Can the field account place some of that energy in one tiny part of space?

In empty space, the answer is yes.

Choose a tiny box. Measure E and B inside it.

The field theory gives an energy amount for that box.

Divide by the box's volume.

The result is **field energy per volume**.

Only after that job is clear do we give it a shorter name: **energy density**, written **u**.

C38.1 — The Vacuum Rule

E contributes one part of the field energy.

B contributes another part.

The two parts add.

In empty space, using SI units:

$$u = \tfrac{1}{2}\varepsilon_0 E^2 + B^2/(2\mu_0)$$

E^2 and B^2 mean that the field amounts are multiplied by themselves.

The symbols ε_0 and μ_0 are vacuum constants.

In SI units, they set the scale of the two parts.

The formula adds the electric and magnetic parts.

C38.2 — What This Does Not Mean

The formula gives energy a local address in this field model.

It does not make energy a fluid or a pile of tiny objects.

A field can store energy without that energy flowing anywhere.

Flow becomes important when the scene changes. A magnet might move, turn, or be released.

C38.3 — The Boundary

This formula is for empty space.

Inside a material, the field and the material respond together.

A material's response can depend on what happened to it earlier.

A real permanent magnet therefore needs a fuller energy rule.

Momentum Through the Gap

C39 — Fields Can Carry Momentum

C39.1 — Energy Crossing a Tiny Window

Imagine a tiny window in empty space.

Fields can move energy through it.

We can ask how much crosses each second through each unit of area.

We can also ask which way the energy goes.

> **Exact note — optional on a first reading**
> The vector **S** records the direction and amount of energy flow. It is called the **Poynting vector**.

C39.2 — Momentum Goes With It

A field can carry momentum as well as energy.

In empty space, the energy flow and field momentum point the same way.

This gives the gap a local momentum account.

> **Exact note — optional on a first reading**
> The vector **g** records field momentum per volume. In empty space:
> $\mathbf{g} = \varepsilon_0 \mathbf{E} \times \mathbf{B} = \mathbf{S}/c^2$
> Here **c** is the speed of light in empty space.

C39.3 — A Picture to Avoid

Field momentum is not a hidden liquid in a pipe.

A magnetic field can store energy even when that energy is not steadily moving away.

Inside matter, the field and the material share the momentum account.

Useful models may divide that account differently. The total must still balance.

C40 — Field Stress

Draw an imaginary closed surface around one magnet.

C40.1 — One Patch

Look at one tiny patch of the surface.

The field account tells us how momentum crosses that patch.

The answer depends on which way the patch faces.

C40.2 — The Whole Surface

Now add the answers from all the patches.

This gives the momentum crossing the whole surface.

If the fields change, the field momentum inside can change too.

The crossing and the inside change together account for the force on the matter.

For a static scene, the inside field momentum does not change.

The surface account is then enough.

> **Exact note — optional on a first reading**
> The rule that gives the answer for every patch direction is the **Maxwell stress tensor**. Adding

> its answers over the closed surface is a **surface integral**. The words name the calculation. They do not mean that empty space is a squeezed solid.

C40.3 — Inside the Magnet

The magnet is solid matter.

The solid must carry the push to the hand.

So its ordinary material stress belongs in the complete account too.

C41 — The Trade Between Field and Matter

We can now follow the push to the hand.

C41.1 — One Tiny Box

Choose a tiny box containing charge and moving charge.

The local fields push on that matter.

The electric part depends on charge in the box.

The magnetic part depends on moving charge there.

> **Exact note — optional on a first reading**
> The force per volume is $\mathbf{f} = \boldsymbol{\rho}\mathbf{E} + \mathbf{J} \times \mathbf{B}$.
> Here $\boldsymbol{\rho}\mathbf{E}$ is the electric part. $\mathbf{J} \times \mathbf{B}$ is the magnetic part.

C41.2 — The Momentum Trade

Suppose the matter gains momentum.

The field account changes by the matching amount.

Momentum has moved between matter and field. It has not disappeared.

C41.3 — The Energy Trade

The field can also give energy to moving charge.

The field's stored energy and energy flow then change.

The total energy account stays balanced.

The magnetic push on one point charge is sideways to its motion.

So that magnetic part does not directly speed up the charge at that instant.

A changing magnetic system can still move energy. Electric fields and moving matter also take part.

C41.4 — A Bar Magnet Has Material Order

A bar magnet is not just loose charge.

Its material has many small magnetic contributions.

Many of them form a lasting local pattern.

The large-scale amount and direction of that pattern is called **magnetization**.

We still need to learn why the pattern forms and lasts.

C41.5 — A Model Must Include the Pattern

A large-scale model needs a way to count the material pattern.

One method treats some of its effects like current tied to the material.

This is not loose charge traveling around an ordinary wire.

> **Exact note — optional on a first reading**
> Current used to represent the material pattern is called **bound current**.

C41.6 — A Model Must Include the Response

The model must also say how the material changes when the field changes.

Different materials can respond differently.

> **Exact note — optional on a first reading**
> A rule for the material's response is called a **constitutive rule**.

C41.7 — Carrying the Push

The fields push on matter at many local places.

Stress in the solid joins those local pushes together.

The magnet's surface then pushes on the hand.

C41.8 — Back to the Hand

The classical chain is:

> field at each place → push on matter → stress through the magnet → push on the hand

The account no longer skips across the gap.

It follows the push from place to place.

But it still does not explain why a permanent magnet stays magnetized.

That question needs atoms and quantum physics.

The Dipole Pattern

C42 — A Simple Faraway Pattern

Return to one bar magnet and the compass.

Move the compass around the magnet. Record the local direction of **B** at many places.

Close to the magnet, small details matter.

Move much farther away than the magnet's length. Many of those details fade from the pattern.

Far away, the main pattern has two lobes.

A small wire loop carrying steady current can have the same faraway shape.

This simple two-lobed field is called a **magnetic dipole field**.

The lobes belong to one field pattern. They are not two isolated magnetic charges hidden at the ends.

They are also not physical objects. They are features of the local field values and their drawing.

The dipole picture works best when the source is small compared with the distance to the probe.

It is an approximation. It does not preserve every detail near a real bar magnet.

C43 — Magnetic Moment

C43.1 — One Vector for Strength and Direction

We need a compact way to compare dipole patterns.

One pattern may be stronger than another. It may also face another way.

A vector can record both facts.

This vector is called the **magnetic moment**. It is written $\boldsymbol{\mu}$.

The amount of μ records the dipole strength. Its direction records which way the pattern faces.

What that strength predicts depends on the test we use.

Next we will use it to predict turning and force.

C43.2 — What the Vector Does Not Mean

μ is part of our dipole model.

It is not a tiny bar magnet hiding inside every object.

A current loop, an atom, and a piece of magnetic material can all have magnetic moments. Their physical origins need not be the same.

We will start with turning.

Orientation and Turning

R-MATH-01 — The Dot Product

Two vectors can point the same way, sideways to each other, or opposite ways.

Sometimes we need one signed number that records this relative direction.

The operation that gives that number is called the **dot product**.

For two vectors **a** and **b**:

> **a · b = amount of a × amount of b × cos θ**

The mark **θ** is the angle between the vectors.

The × marks on the right mean ordinary multiplication of numbers. They are not vector cross products.

The word **cosine** names the angle factor.

- The factor is **1** when the vectors point the same way.
- It is **0** when they meet at a right angle.
- It is **−1** when they point in opposite ways.

A dot product is a signed number. It is not a new vector or physical object.

In plain words, it answers one question:

> How much do these two vectors point together rather than against each other?

C44 — Orientation Energy

Place an ideal small dipole in an external magnetic field **B**.

Hold B steady with outside equipment.

This lets us study only the dipole's turning.

Physicists call a field held this way a **prescribed external field**.

Also keep the amount of μ fixed.

The turning-energy rule is:

> **$U_{mech} = -\mu \cdot B + \text{constant}$**

When μ points along B, the energy is lower.

When μ points sideways to B, it lies between the other two cases.

When μ points against B, the energy is higher.

The minus sign in the formula creates that result.

The added constant only chooses where we call the energy zero.

This is not the total energy of the dipole, field, and equipment. It is the part we need for the dipole's orientation.

C45 — Dipole Torque

A dipole that does not point along B tends to turn.

The torque is:

> $$\boldsymbol{\tau} = \boldsymbol{\mu} \times \mathbf{B}$$

The cross product gives the turning direction and amount.

The turn tends to bring μ into line with B.

When μ and B are aligned, the torque is zero.

When they point exactly opposite ways, the torque is also zero.

But that balance is easy to upset. A tiny turn makes the dipole keep turning away from that position.

Torque turns. It does not by itself push the dipole's center one way.

C46 — Turn Without a Net Pull

Suppose B has the same amount and direction across the whole small dipole.

This is a **uniform field**.

The dipole may turn if μ does not point along B.

But an ideal point dipole in this uniform field has no net push on its center.

Different parts can still feel forces. In the ideal uniform case, those pushes cancel.

So a field can turn a dipole without pulling its center one way.

From One Dipole to Two Magnets

C47 — A Field Gradient Can Pull

Let B be stronger on one side of a dipole than on the other.

The two sides now sit in different field strengths.

Their pushes may no longer cancel.

So a change in B across space can push the dipole's center.

That change across space is called a **field gradient**.

This does not mean every dipole moves toward the strongest B.

The result also depends on the dipole's direction and on how B changes.

C48 — The Point-Dipole Force

C48.1 — The Small, Slow Scene

Imagine an ideal dipole so small that we treat it as one point.

Keep its magnetic moment μ fixed.

Let equipment outside the dipole provide B.

Move the dipole slowly enough that we can ignore waves during this step.

The next rule belongs only to this scene.

C48.2 — Compare Nearby Places

At one place, μ may fit B better than at another.

A better same-direction fit means a lower orientation energy.

So the dipole is pushed toward the direction where that fit improves fastest.

This is the same lower-energy rule we used earlier.

Exact note — optional on a first reading
The exact point-dipole force is $\mathbf{F} = \nabla(\boldsymbol{\mu} \cdot \mathbf{B})$. First find the dot product $\boldsymbol{\mu} \cdot \mathbf{B}$. The mark ∇ then points toward the direction where that number rises fastest.

C48.3 — Where the Shortcut Fails

The rule assumes:

- a fixed μ;
- an outside field B;
- a dipole small enough to treat as a point;
- motion slow enough to ignore wave effects here.

Nearby real magnets can change each other's state.

Then we need the fuller energy or stress account.

C49 — Two Magnets Form One Arrangement

C49.1 — Each Magnet Takes Part

With two magnets, each helps create the electromagnetic scene.

Each also responds to that scene.

Their positions and directions belong to one connected problem.

C49.2 — A Faraway Shortcut

Far apart, we can treat each magnet as a dipole.

As a shortcut, first calculate the field from one magnet.

Then calculate how the other responds there.

It does not mean that one magnet acts while the other does nothing.

C50 — Why a Flip Can Reverse the Force

C50.1 — Flip One Magnet

Hold the two bar magnets end to end.

Now flip one magnet.

Its magnetic moment reverses direction.

The dot products change. The energy at each position changes too.

So the force can reverse.

A pull can become a push, or a push can become a pull.

C50.2 — The Shape Still Matters

That result belongs to this end-to-end setup.

Side-by-side magnets can behave differently.

Nothing traveled along a field-line string.

The flip changed the local field, energy, and stress pattern.

C51 — Real Magnets Fill Space

C51.1 — Many Local Pushes

A real bar magnet is not a point.

Its material fills a volume.

The field and the local push can differ from one part to another.

Those many pushes make stress inside the solid.

The solid carries the stress to its surface and then to the hand.

We can measure the total force.

We can also measure how that force is spread across the surface.

C51.2 — When the Dipole Picture Works

The point-dipole picture works best when the distance is large compared with the magnet's size.

It can fail when magnets are close.

It can also fail when their magnetization changes strongly.

Different models may divide the inside account in different ways.

They must still agree on the total force and the push at the surface.

C51.3 — What It Still Cannot Explain

The dipole picture explains useful direction and distance patterns.

It does not explain why the material keeps a magnetic moment.

That question remains open for the next part of the book.

Current, Flux, and Circulation

C52 — Electric Current

C52.1 — Choose Which Way Counts as Positive

Imagine a flat cut across the wire loop.

Draw an arrow through the cut.

The arrow says which crossing direction counts as positive.

That choice is the surface's **orientation**.

Positive charge crossing with the arrow adds to the count.

Negative charge crossing with the arrow subtracts.

Negative charge moving the other way therefore adds positive current.

C52.2 — Count Charge Crossing Each Second

The signed charge crossing each second is the **electric current**.

We write it as **I**.

If I is changing, we may want its value at one exact instant.

> **Exact note — optional on a first reading**
> The instant-by-instant rate is **I = dQ_cross/dt**. The mark **d/dt** means “how fast is this count changing now?”

One ampere means one coulomb of charge per second.

C52.3 — One Total and Many Local Parts

I is one number for the whole cut.

But the current may not be spread evenly.

The vector **J** describes the current near each small place.

Its direction shows the local flow direction.

Its amount tells the current passing through each unit of area there.

We call J the **current density**.

R-EM-02 — Two Ways to Add Local Field Values

Sometimes we need one answer for a whole surface or path.

We will build each kind of answer separately.

R-EM-02.1 — Through a Surface

Choose a surface.

Choose which of its two sides counts as positive.

Now divide the surface into tiny patches.

At each patch, keep only the part of the field pointing through it.

A sideways part counts as zero for that patch.

Multiply the through-part by the patch's area.

Then add the signed answers from every patch.

This total is called the field's **flux** through the surface.

R-EM-02.2 — The Flux Shorthand

> **Exact note — optional on a first reading**
> For any vector field G: **flux of G = ∫ G · n dA**. The arrow **n** points toward the chosen positive side. The mark **dA** means one tiny patch. The integral sign means "add all the patches."

Reverse the chosen side and the flux changes sign.

Flux is a mathematical total. It does not always mean that matter flows through the surface.

R-EM-02.3 — Current Is One Use of Flux

J tells current near each patch.

Adding its through-parts gives the total current I.

> **Exact note — optional on a first reading**
> The exact sum is **I = ∫ J · n dA**.

This is why total current and current density are not the same thing.

R-EM-02.4 — Around a Path

Now choose a closed path and a direction around it.

Break the path into tiny pieces.

At each piece, keep the field part pointing along the path.

Multiply by that piece's length.

Then add all the answers around the loop.

This total is called **circulation**.

R-EM-02.5 — The Circulation Shorthand

> **Exact note — optional on a first reading**
> For any vector field G: **circulation of G =** $\oint \mathbf{G} \cdot d\boldsymbol{\ell}$.
> The mark $d\boldsymbol{\ell}$ is one tiny directed path piece. The circle on the integral sign says that the path closes.

Reverse the path direction and the circulation changes sign.

Circulation is a mathematical total. It does not require a fluid to turn around the path.

C53 — Making the Two Totals Local

Flux and circulation describe whole surfaces and paths.

Now shrink those shapes around one point.

C53.1 — A Tiny Closed Box

Put a tiny closed box around the point.

Find the outward flux through the box.

Divide by the box's volume.

Make the box smaller and watch which value the answer approaches.

That local number is the field's **divergence**.

> **Exact note — optional on a first reading**
> Divergence is written $\nabla \cdot \mathbf{G}$.

Positive divergence means a local outward pattern.

Negative divergence means a local inward pattern.

Divergence is one signed number.

C53.2 — A Tiny Loop

Draw a tiny loop around the point.

Find the circulation around it.

Divide by the loop's area.

Shrink the loop and repeat with the loop facing different ways.

The resulting local direction and amount form the field's **curl**.

> **Exact note — optional on a first reading**
> Curl is written $\nabla \times \mathbf{G}$.

Its direction follows the right-hand rule.

C53.3 — The Short Difference

Divergence uses a tiny closed box. It asks about a local outward pattern.

Curl uses a tiny loop. It asks about a local circling pattern.

Neither one says that the field is a material fluid.

Maxwell's Four Local Equations

C54 — Steady Current Makes B Circle

Send a steady current through one straight part of the wire loop.

Move the compass around the wire.

Its arrow makes a circling pattern around the current.

Curl is our local measure of that pattern.

When nothing changes with time, local current is tied to local curl in B.

> **Exact note — optional on a first reading**
> The steady rule is $\nabla \times \mathbf{B} = \mu_0 \mathbf{J}$. The left side is curl in B. The right side contains local current density J. The constant μ_0 sets the SI scale.

This is only the steady case.

A changing electric field adds one more part.

C55 — Changing B Makes E Circle

C55.1 — The Broken Wire Loop

Take the wire loop and leave a small break in it.

Now change the magnetic field through the loop.

An electric field can circle all the way around, even across the break.

The wire is not creating that field.

It only supplies charges that can respond.

C55.2 — The Whole-Loop Job

Choose a positive direction around the loop.

Also choose the matching positive side of the loop's inside surface.

The electric circulation around the loop opposes the change in magnetic flux through it.

This is called **Faraday's law**.

> **Exact note — optional on a first reading**
> For a fixed loop and surface: $\oint \mathbf{E} \cdot d\boldsymbol{\ell} = -d/dt \int \mathbf{B} \cdot \mathbf{n}\, dA$.
> The left side adds E around the loop. The right

side watches magnetic flux. The minus sign sets the opposing direction. The right-hand rule links the two chosen directions.

C55.3 — The Job at One Place

The same physical link can be stated at one place.

A changing B is tied to local curl in E.

Exact note — optional on a first reading
The local rule is $\nabla \times \mathbf{E} = -\partial \mathbf{B}/\partial t$. If B stops changing, this changing-B part is zero. E itself need not be zero.

C56 — Changing E Helps Make B Circle

C56.1 — The Gap Between Two Plates

Imagine two conducting plates with an empty gap between them.

Charge reaches one plate and leaves the other.

No charged matter crosses the gap.

Even so, E in the gap changes.

B can circle around that changing electric region.

C56.2 — The Missing Job

The full B rule must therefore count changing E as well as current.

Changing E can help make B curl where no charged matter crosses.

> **Exact note — optional on a first reading**
> The added part is $\mu_0\varepsilon_0\,\partial E/\partial t$. The quantity $\varepsilon_0\,\partial E/\partial t$ has the same units as current density J. It is called the **displacement-current term**. The name does not mean that charge crosses the gap.

C57 — Four Jobs Before Four Equations

C57.1 — The First-Reading Version

- Charge is tied to the outward pattern of E.
- B has no isolated starting or ending point.
- Changing B makes E circle.
- Current and changing E make B circle.

These rules connect nearby places and times.

They do not choose one complete field scene by themselves.

The exact lines below are here for readers who want the mathematical form now.

They are not needed to continue the first reading.

C57.2 — Exact Line One: Charge and E

> **Exact note — optional on a first reading**
> $\nabla \cdot \mathbf{E} = \rho/\varepsilon_0$
> The left side asks for divergence in E. The right side contains charge per volume, written ρ. The constant ε_0 sets the SI scale.

C57.3 — Exact Line Two: No Isolated Ends of B

Exact note — optional on a first reading

$\nabla \cdot \mathbf{B} = 0$

The local outward total for B is zero. This does not say that B is zero. The ends of a bar magnet are not separate magnetic charges.

C57.4 — Exact Line Three: Changing B and E

Exact note — optional on a first reading

$\nabla \times \mathbf{E} = -\partial \mathbf{B}/\partial t$

The left side asks for curl in E. The right side asks how B changes there. The minus sign sets the opposing direction.

C57.5 — Exact Line Four: Current, Changing E, and B

Exact note — optional on a first reading

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \varepsilon_0 \, \partial \mathbf{E}/\partial t$$

The two parts on the right are local current density and changing E. Both help make curl in B.

C57.6 — Rules Shared by All Four Lines

Exact note — optional on a first reading

Use one place, one time, and one reference frame in each line. The mark **$\partial/\partial t$** asks how fast a field changes at one fixed place. E uses newtons per coulomb. B uses tesla. ρ uses coulombs per cubic meter. J uses amperes per square meter.

These equations count all charge and current represented at each tiny place.

A material model may group some effects in another way.

We must not silently mix the two accounts.

C57.7 — Charge Does Not Vanish

Charge can move into or out of a tiny box.

It cannot simply disappear from the account.

If more current leaves than enters, the charge inside goes down.

> **Exact note — optional on a first reading**
> The local charge rule is $\partial\rho/\partial t + \nabla \cdot \mathbf{J} = 0$.

Selecting a Solution and Making a Wave

R-EM-03 — A Rule Is Not Yet One Solution

Maxwell's equations tell us what every allowed local field must obey.

Many different field patterns can obey the same local equations.

R-EM-03.1 — Choose the Scene

First say what the fields are when the problem starts.

Those starting values are called **initial data**.

Next say what happens at the edges of the region.

Those edge rules are called **boundary data**.

The shapes and positions of the wires, magnets, and empty spaces matter too.

R-EM-03.2 — Matter Can Respond

A material can shift positive and negative charge slightly apart.

We need a local amount and direction for that separation.

It is called **polarization**, written **P**.

P summarizes electric separation inside the material.

A material can also contain many small magnetic moments.

Their magnetic moment per volume is **magnetization**, written **M**.

M summarizes the material's magnetic order.

R-EM-03.3 — Two Helper Fields

Two helper fields are often used in large-scale matter equations. They are called **D** and **H**.

D and H help us choose which effects to count as sources and which to count as material response.

That split is a bookkeeping choice.

They are bookkeeping helpers. They are not deeper replacements for E and B.

R-EM-03.4 — The Material Needs Its Own Rule

We still need to know how P and M respond to the fields.

The answer may depend on field strength. It may also depend on what happened to the material earlier.

A permanent magnet keeps some of that history.

A rule that describes this response is called a **constitutive rule**.

So Maxwell's equations are only one part of a real magnet problem.

We also need the starting scene, the edges and shapes, and a rule for the material.

C58 — Electromagnetic Waves

Look at empty space with no local charge or current.

A changing E can make B circle.

A changing B can make E circle.

Together, those changes can move from one region to the next.

This is an **electromagnetic wave**.

In empty space, it moves at the speed of light.

We write that speed as **c**.

> **Exact note — optional on a first reading**
> The exact relation is $\mathbf{c = 1/\sqrt{(\mu_0 \varepsilon_0)}}$. The constants ε_0 and μ_0 set the electric and magnetic scales in empty space. The square-root mark asks for the number that gives $\mu_0\varepsilon_0$ when multiplied by itself.

In a simple light wave, E and B point sideways to the direction of travel. They also point at right angles to each other.

The energy flow points in the travel direction. The wave also carries momentum.

The field disturbance propagates at a finite speed. A change here does not update the distant field instantly.

A static resting magnet does not continually send out such radiation.

This is the classical wave account. We have not yet introduced photons or a quantum field.

Reference Frames and One Electromagnetic Field

C59 — A Full Reference Frame

Earlier, a reference frame told us what counted as rest or motion.

Now we need its full job.

Imagine a grid of rulers and clocks.

An **event** means one happening at one place and time.

The grid gives each event a place and a time.

It also tells us how fast something moves through the grid.

That whole measuring system is a **reference frame**.

A frame that is not speeding up, slowing down, or turning is an **inertial frame**.

Two such frames can move steadily past each other.

They may give different numbers for place, time, speed, charge density, current, E, and B.

These are not mere changes in how something looks.

Instruments moving with each frame measure that frame's values.

Relativity gives rules for changing the full set of numbers from one frame to the other.

When transformed together, both descriptions predict the same outcome for the experiment.

Never combine E from one frame with v from another frame in the Lorentz-force law.

C60 — One Field, Two Frame-Dependent Parts

Suppose one inertial observer measures an electric field and no magnetic field at one event.

An observer moving past the first may measure both E and B at that same event.

The reverse can also occur. A field that is purely magnetic in one frame can include an electric part in another.

The moving observer did not create a new field just by moving.

E and B are frame-dependent parts of one electromagnetic field.

Changing frame can mix the two parts. Special relativity gives the exact mixing rule.

In any one frame, use the E and B measured in that frame to predict force.

This does not mean every magnetic field can be removed by changing frame.

We need two tests before we can say when that is possible.

Invariants and the Classical Boundary

C61 — Two Checks Shared by Every Frame

Changing inertial frame can mix E and B.

Sometimes one part can become zero. Sometimes neither can.

Two checks tell us which kind of case we have.

C61.1 — Do E and B Share a Direction?

The first check asks whether E and B point partly together.

The answer is zero when they are at right angles.

It is also zero when either field is zero.

If the answer is not zero, no inertial frame can make either field vanish at that event.

> **Exact note — optional on a first reading**
> This first check is the dot product $\mathbf{E} \cdot \mathbf{B}$.

C61.2 — Which Part Wins the Second Check?

The second check compares the amount of E with the amount of B.

We first multiply B by the speed of light so the two amounts can be compared.

The answer can favor the electric part, favor the magnetic part, or balance.

> **Exact note — optional on a first reading**
> The second check is $\mathbf{c^2B^2 - E^2}$. A negative answer is called **electric-dominant**. A positive answer is called **magnetic-dominant**. Zero means this check is balanced.

C61.3 — Why the Checks Matter

Different inertial observers may measure different E and B parts.

Yet every inertial observer gets the same two check results at the same event.

Such shared results are called **invariants**.

Their job here is narrow.

They tell us whether a change of frame can remove E or B.

R-REL-01 — Can One Part Disappear?

R-REL-01.1 — When B Can Disappear

First, E and B must not share a direction.

Second, the electric side must win the strength check.

When both facts hold, some inertial frame measures B as zero at that event.

> **Exact note — optional on a first reading**
> For a nonzero field, $\mathbf{B} = \mathbf{0}$ in some inertial frame exactly when $\mathbf{E} \cdot \mathbf{B} = 0$ and $c^2B^2 - E^2 < 0$.

R-REL-01.2 — When E Can Disappear

Again, E and B must not share a direction.

This time, the magnetic side must win the strength check.

When both facts hold, some inertial frame measures E as zero at that event.

Exact note — optional on a first reading
For a nonzero field, $\mathbf{E} = \mathbf{0}$ in some inertial frame exactly when $\mathbf{E} \cdot \mathbf{B} = 0$ and $c^2B^2 - E^2 > 0$.

R-REL-01.3 — When Neither Can Disappear

There is one more important case.

The field is not zero, but both checks give zero.

Then no inertial frame can make the field purely electric or purely magnetic.

A plane electromagnetic wave is an example.

Exact note — optional on a first reading
A nonzero field with $\mathbf{E} \cdot \mathbf{B} = 0$ and $c^2B^2 - E^2 = 0$ is called a **null field**.

R-REL-01.4 — One Place and Time

Each check concerns one event: one place at one time.

A single frame may not remove the same field part everywhere in an uneven or changing scene.

So the claim “magnetism can always be transformed away” is false.

C62 — What Relativity Explains

Special relativity explains why changing inertial frame can mix E and B.

It shows that they are two frame-dependent parts of one electromagnetic field.

That explains a deep link between electricity and magnetism.

It does not explain why matter contains tiny magnetic parts.

It also does not explain why many line up and stay that way.

Those questions need quantum physics and the history of the material.

C63 — The Job of Classical Electromagnetism

Give the theory the charges, currents, starting scene, shapes, edges, and material rules.

It can then predict:

- large-scale E and B fields;
- pushes and turns;
- energy, momentum, and stress;
- light and other electromagnetic waves.

For large-scale problems that fit the theory, these predictions can be extremely accurate.

But charge, current, and material response enter as starting facts.

Classical electromagnetism does not explain why a permanent magnet's tiny magnetic parts form a lasting order.

That is the next question.

The Classical Account and Its Boundary

We began with two bar magnets and two hands.

The hands feel contact pushes at the grips. Those pushes are the last links in a chain. Stress passes through each magnet, into the skin, and through the body.

The chain also crosses the gap.

At every place, electric and magnetic fields have measurable amounts and directions. They predict how charge will move. Their patterns depend on where the magnets are and which way they face.

The fields do more than label the empty space. They can store energy, carry momentum, and exert stress. Matter and fields exchange momentum locally. A free magnet can move or turn. A held magnet passes the load through its solid body to the hand.

Maxwell's equations connect charge, current, and changing fields from one place to the next. Special relativity shows that electric and magnetic fields are two frame-dependent parts of one electromagnetic field.

So the classical account gives a real answer to what is happening between the magnets. The gap has a measurable physical state. That state follows local rules and carries energy and momentum. Together, the fields and matter account for the push or pull that reaches the hands.

But the account begins with a magnet that already has lasting magnetic order. Classical electromagnetism can use that material fact. It does not explain why the order exists or why it remains after the preparing field is gone.

The next chapter begins inside the magnet, with atoms and electrons. It asks how lasting magnetic order is made.

Why This Solid Can Be a Magnet

C64 — What Is Inside a Magnet?

A permanent magnet can keep its effect after the preparing field is gone.

Classical field rules can use that fact. They do not explain it.

We must now look inside the material.

Picture the same piece before and after magnetic preparation. Before preparation, it may have little lasting outside effect. After preparation, part of its inside pattern can remain changed.

A solid is made of **atoms**. Atoms contain electrons.

An **electron** is a tiny part of matter with negative electric charge. Electrons help hold atoms together. They also help join atoms into a solid.

The lasting magnetism of an ordinary bar magnet comes mainly from its electrons. The nuclei can have magnetic effects too. Those effects are not the main source here.

This does not mean that each electron is a tiny bar magnet. Many electron effects cancel. Some electrons take part in shared patterns that spread through the material.

The whole arrangement matters.

We need one more idea before we can discuss that arrangement. We need to know what a quantum rule can tell us about one electron.

C65 — What Can a Repeated Test Tell Us?

Take a large set of electrons. Prepare each one in the same way.

Choose one test. Let a detector leave one clear record for each electron.

The first record might be one result. The second might be another. After many runs, a pattern appears. Some results happen more often than others.

Quantum theory gives a rule for this pattern.

The rule tells us which records are possible for the chosen test. It also gives the chance of each possible record.

Now choose a different test. The same rule must give the possible records and chances for that test too.

This complete prediction rule is called a **quantum state**.

A quantum state belongs to the prepared situation. It is not one detector record. It is not an unseen card holding a fixed answer for every test.

The detector records what happened in one run. The state lets us predict the pattern across many runs prepared the same way.

When several electrons act together, quantum theory uses one state for the group. That group state gives the possible records and chances for tests on the group.

We need group states because a solid contains many electrons. The lasting magnetic pattern is a group result.

Before we build that group, we need one measured property of a single electron.

C66 — Why Does One Electron Have a Magnetic Effect?

Prepare one electron the same way many times.

Choose a direction for the test. The direction belongs to the test. It might point up from the table.

Run the test once. The detector can leave one of two records along that chosen direction.

Repeat the test with newly prepared electrons. The state gives the chance of each record.

This two-result measured property is called **spin**.

The name can be misleading. An electron is not a tiny ball with a known surface turning around.

Spin is a quantum property. It tells us how the electron can answer a spin test. It also helps set how the electron responds to a magnetic field.

Change the test direction, and the chances can change. Quantum theory does not replace this rule with a hidden list of ordinary arrows pointing in space.

So far we have one electron, a state, and a measured property.

We still need to connect that property to magnetism.

C67 — How Does Spin Add a Magnetic Contribution?

Put the electron in a magnetic field.

Choose the spin-test direction along the field.

The two possible spin results can now have different energy numbers.

Change the field, and those energy numbers can change. This means spin helps set the energy of an electron in a magnetic field.

That energy link is spin's **magnetic contribution**.

One electron's contribution is tiny. In a solid, many contributions can add. Many can also cancel.

Spin is not the only electronic source. The quantum pattern of an electron inside an atom can contribute too.

The balance depends on the material.

A bar magnet contains a vast number of electrons. The next question is how their allowed group states can favor a shared pattern.

How Many Small Effects Become One Lasting Magnet

C68 — What Does “Same Kind of Electron” Mean?

A solid contains many electrons of the same kind.

Electrons of the same kind have the same built-in properties. They have the same electric charge and the same mass.

Imagine placing a hidden name tag on each electron. Now switch the tags.

No test result changes. The tags were never part of the physical arrangement.

Physicists say the electrons are **identical**. Here, identical means that hidden name tags cannot tell them apart.

Their shared quantum state must respect that fact.

This does not mean two electrons must have the same test result. It means the group rule cannot treat a switched name tag as a new physical case.

That fact helps limit the group arrangements.

C69 — Which Group Arrangements Are Allowed?

Prepare two electrons in the same atom.

The state for one electron includes spin and other parts. The other parts include how the electron is spread through the atom.

First imagine that every part of the two electron states matches. They have the same spread through the atom. They also have the same spin result along one chosen test direction.

That exact duplicate is not an allowed group arrangement.

Now keep the same spread but use opposite spin results along that test direction. That arrangement can be allowed.

If the ways they spread differ, two electrons can also have the same spin result.

The rule removes the exact duplicate. It does not ban every same-spin pair or every close pair.

This limit is called the **Pauli exclusion rule**.

The Pauli rule is not a new force. It does not push electrons apart like two magnets with matching ends.

It is a rule about which group states are allowed.

We now have fewer allowed arrangements. The next step is to compare their energy numbers.

C70 — Why Can One Arrangement Cost Less?

Take one material. Focus on two nearby parts that can contribute to magnetism.

A material model may compare two allowed arrangements.

Use the same chosen test direction for both nearby parts. In one arrangement, repeated test records tend to match. In the other, they tend to differ.

Here, same-way names that record pattern. It does not mean each electron carries a hidden little arrow through space.

Both arrangements must first obey the group-state rules.

The calculation also includes the electric interaction among the charged parts.

When both ingredients are used, the two allowed arrangements can have different energy numbers.

The relevant part of this energy difference is called **exchange energy**.

Exchange is not another force added to electricity. It is part of the energy comparison between allowed quantum arrangements.

The answer depends on the material and the model. In some materials, the same-way pattern has lower energy. In others, an opposite pattern or no lasting order is favored.

So the Pauli rule alone does not align anything. The useful chain is longer:

1. The group-state rule limits the allowed arrangements.
2. The electromagnetic interaction helps set their energy numbers.
3. The material model tells us which nearby pattern costs less.

That lower-energy choice can be repeated across part of a solid.

Quick Refresh — From One Electron to a Material Choice

A quantum state predicts possible test records and their chances.

Spin is a two-result measured property that helps set an electron's magnetic response.

A group of identical electrons cannot use every imagined arrangement.

The allowed arrangements can have different energy numbers. Which arrangement costs less depends on the material.

C71 — How Can a Nearby Pattern Spread?

Return to the material where a same-way nearby pattern has lower energy.

One nearby pair is not a bar magnet.

Now add more magnetic parts. Each part has nearby partners. The material's energy rule applies throughout the solid.

Compare two arrangements of one small region.

In the first, many neighboring parts use the lower-cost pattern. In the second, more neighbors use a higher-cost pattern.

The first arrangement can then have a lower total energy number. Other parts of a real material model can matter too.

This energy difference can favor one linked pattern across many neighboring parts.

This shared pattern is called **magnetic order**.

Magnetic order does not mean perfect stillness. The parts remain quantum objects. Heat also keeps disturbing the pattern.

Order lasts only when the energy preference is strong enough under the present conditions.

Raise the temperature far enough, and heat can break the long-range pattern. Cool the material again, and order may return.

Not every solid has the needed energy rule. Some favor opposite nearby patterns. Some have strong cancellation. Some lose order at ordinary temperatures.

A material with local magnetic order has passed one test. It still may not act like a strong bar magnet outside the solid.

C72 — Why Can Ordered Material Look Weak?

Picture a fresh piece of ordered magnetic material.

One part of the piece may favor one direction. Another part may favor a different direction.

Inside each part, many small magnetic contributions share a pattern.

Such a region is called a **magnetic domain**.

A thin changeover region lies between two domains. This is called a **domain wall**.

Domains can point in different directions. Their large-scale effects can then cancel.

The material may have strong magnetic order inside each domain while showing little outside effect as a whole.

This solves an important puzzle. Local order does not guarantee a strong permanent magnet.

The sizes and directions of the domains matter too.

A preparing magnetic field can change that domain pattern.

C73 — Why Does the New Pattern Stay?

Place the material in a strong outside magnetic field.

Domains that fit the field can grow. Some local magnetic patterns can turn. Domain walls can move.

The material now has more magnetic contribution in one shared direction.

Remove the preparing field.

The material does not always return to its earlier domain pattern.

Why not?

Some directions cost less than others inside the solid. To reverse a region, the material may first have to pass through a higher-energy arrangement.

That higher part of the route is an **energy barrier**.

Real solids also contain grain edges, mixed atoms, gaps, and other defects. A domain wall can catch on them. The defect can block a change that would otherwise continue.

The barrier and the blocked walls give the material memory.

The magnetic effect left after the preparing field is removed is called **remanence**.

Remanence is not eternal. A strong opposite field can change it. A hard strike can sometimes help changes begin. Enough heat can also help the material cross barriers.

A permanent magnet is therefore a prepared material that resists reversal under ordinary use.

C74 — Why Aren't Most Objects Strong Magnets?

A strong permanent magnet needs several links to work together.

The material needs electronic magnetic contributions that do not all cancel at once.

Its allowed many-electron arrangements must favor useful local order.

That order must survive at the working temperature.

The prepared state must leave enough magnetic contributions uncancelled.

Energy barriers must then slow reversal. In many materials with several domains, defects that pin domain walls help too.

Miss one link, and the object may show little lasting outside field.

A paper clip can respond to a nearby magnet. It does not have to be a strong permanent magnet of its own. The nearby field changes its magnetic pattern. Much of that change can fade when the magnet leaves.

A finished bar magnet has a different history. It was made from a suitable material and put through a preparing process. Part of the prepared pattern remains.

That remaining pattern is the material source used by the large-scale field account.

Quick Refresh — How the Magnet Keeps Its Effect

Electrons obey group-state rules. The allowed arrangements can have different energy numbers.

In a suitable material, the lower-cost choices support local magnetic order.

The ordered material can split into domains. Different domains may cancel one another.

Preparation changes the domain pattern. Barriers and blocked domain walls keep part of that change from reversing.

The result is a lasting, uncancelled magnetic effect.

From Internal Order to the Pull on the Hand

C75 — How Do We Summarize a Vast Number of Electrons?

A bar magnet contains too many electrons to track one at a time.

For a large-scale field problem, we do not need every electron record. We need a useful summary of the prepared material.

Choose one small piece of the magnet. The piece must be large enough to contain many atoms. It must also be small compared with the whole bar.

Combine the magnetic contributions inside that piece. Keep their amount and shared direction.

Now divide by the piece's volume. The volume is the amount of space the piece fills.

This large-scale summary is called **magnetization**.

Magnetization can change from place to place inside the same magnet. A domain has a fairly steady local direction. A domain wall has a changing direction.

Before preparation, different domains may cancel. The field outside the whole piece can then be small.

After preparation, more domains support one overall direction. The uncancelled magnetization changes the field map outside the piece.

Magnetization is not the same thing as the magnetic field B.

Magnetization describes the material's magnetic contribution. B describes the local field used in the force rule taught earlier.

A material model connects them. The exact bookkeeping can change with the problem and the units being used.

For our main question, the job is simple. The prepared electron and domain pattern gives the magnetization. Classical field rules can use that magnetization as material input.

Now the inside story can join the earlier field story.

C76 — How Does the Load Reach the Hand?

Return to the two bar magnets held in two hands.

Each magnet has a prepared magnetization pattern. That pattern helps set the magnetic field around and inside the magnet. Any electric field has other sources, such as electric charge.

Bring the magnets closer.

In an ordinary held-magnet test, each magnet's prepared inside pattern may change very little. Even so, the shared field pattern changes. A strong enough nearby field can also change the inside pattern.

The two magnets now form one electromagnetic arrangement. Each helps set the fields. Each also responds to the fields in the shared arrangement.

The field is not one object sent out by the first magnet and caught by the second. The field has a local value at every place in the model.

The local values change across the gap and through the solids.

Inside each magnet, charged matter and the material's magnetic parts respond to those local fields. The tiny pushes are spread through the material. They do not all act at one point.

Those pushes make stress in the solid magnet. The solid carries the load to the place where the hand grips it.

The magnet presses on the skin. The skin changes shape. Muscles and joints help hold the magnet. Nerves send signals to the brain.

That is how the person feels the magnetic push or pull.

The hand does not reach across the gap. It feels the magnet where the hand and magnet touch.

If the hand lets go, the same unbalanced load can make the magnet move or turn.

C77 — What Does the Large-Scale Answer Explain?

We can now join the large-scale chain.

Start with two prepared magnets. Their material has lasting magnetization.

Give classical electromagnetism the magnet shapes, positions, magnetization, and other needed material rules.

The theory gives local electric and magnetic field values.

Those field values predict local pushes on charged matter.

The same field account tracks energy, momentum, and stress in the gap.

The material carries the distributed load through each solid.

At the grip, the magnet presses on the hand. The body turns that contact load into the feeling of a push or pull.

This is a complete classical answer to the reporter's first question. It tells us how to calculate the large-scale scene and follow the load to the hand.

The quantum material story supplies the fact that classical theory had to take as input. It explains why a suitable prepared solid can keep a magnetization pattern.

The two levels do different jobs.

The material model explains the lasting order inside the magnet.

The classical field model explains the large-scale fields, force, energy, momentum, and stress.

Neither level needs field lines to be physical strings. Empty space does not become a hidden solid. The field and stress are parts of a tested physical account.

We can stop here for many magnet problems.

The reporter also asked for the deepest available account. For that, the electromagnetic field must follow quantum rules too.

What the Deeper Quantum Theory Adds

C78 — How Can a Field Follow Quantum Rules?

Earlier, we built a field map.

At each place and time, the map gives electric and magnetic field values. That classical map works extremely well for many large-scale problems.

Quantum theory asks for a different kind of description.

Prepare the electromagnetic field and the nearby matter in the same way many times. Choose a test. Let a detector leave a record in each run.

The theory must give the possible records and their chances. It must also connect records made at different places or times.

The prediction rule is a quantum state. In this description, the electromagnetic field is a **quantum field**.

A quantum field is not a classical field map with random shaking added.

The quantum state and the chosen test work together to set the possible records and chances. One run gives actual records. Repeated runs reveal the predicted pattern.

Under suitable conditions, repeated measurements can have averages that follow the classical field rules very closely. That is why a classical field map remains useful for the bar magnets in our hands.

A stable average across many runs does not mean every single run carried that same exact classical field value.

The quantum description does not switch on only because something is small. It is the deeper rule at every scale. The classical description is an excellent large-scale approximation for the right states and tests.

This quantum field can carry energy and momentum. In radiation tests, detectors can record separate energy transfers.

C79 — Do Held Magnets Send Out Radiation?

Shine a very weak light at a detector.

The detector does not always receive energy as one smooth smear. It can record separate transfers of electromagnetic energy.

One detectable unit of that radiation is called a **photon**.

A photon can transfer energy and momentum to matter. A source can emit one. A detector can absorb one.

The word packet can help with the energy transfer. It must not turn the photon into a hard little pellet following a known classical path.

The quantum state may spread across a large region. The detector leaves one local record. Quantum theory predicts the chances for where and how records occur.

Now return to two magnets held still.

Their steady magnetic interaction does not require a stream of real photons leaving one magnet and striking the other.

Real outgoing photons carry radiation. A truly steady electromagnetic arrangement is not continually losing energy as outgoing radiation.

Real magnets at ordinary temperatures can also give off heat radiation. That separate fact does not make photon rain the cause of the steady magnetic pull.

Move or shake a charge, switch a field, or disturb a magnet quickly, and radiation may be produced. That changing case is different from the held, nearly static case.

So real photons matter to light and radiation. They are not the missing tiny messengers between resting magnets.

C80 — Are Tiny Messengers Crossing the Gap?

People often say that magnets push by “exchanging virtual photons.”

That sentence can create the wrong picture.

Some deeper quantum problems are too hard to solve in one step. Physicists break the calculation into an organized series of simpler terms.

Some internal terms in that calculation are called **virtual photons**.

A virtual photon is not a detector record. It is not a little object seen leaving one magnet, crossing the gap, and arriving at the other.

It is an internal part of the calculation, so it should not be given a measured path or travel story.

The phrase “borrowed energy” is also misleading. No small messenger breaks the energy rule for a moment and repays it later.

Energy remains accounted for in the complete physical process.

Virtual-photon language can be useful to a person doing the calculation. It is not the ordinary substance between the magnets.

For our static magnets, the useful large-scale account remains the local field, energy, momentum, and stress account already built.

C81 — What Joins Charged Matter and the Field?

We now have two quantum pieces.

Electrons follow quantum rules. The electromagnetic field follows quantum rules too.

We need one theory that tells us how charged matter and that field act together.

That theory is **quantum electrodynamics**, usually shortened to **QED**.

QED predicts possible records and their chances for electromagnetic processes. It covers light emission, light absorption, and many interactions involving electrons and other charged particles.

It also predicts the electron's magnetic response. Careful measurements agree very closely.

QED has more than one allowed mathematical description of the same physical scene. Parts that change between those descriptions cannot be treated as separately measured objects. The predicted detector records stay the same.

This fact helps explain why an internal calculation term is not automatically a thing traveling through space.

QED starts with tested rules and measured inputs, including the electron's charge and mass. It uses those inputs. It does not explain why nature chose all of them.

QED gives the deeper electromagnetic rules. It does not build a finished bar magnet in one jump.

C82 — Can One Theory Build a Bar Magnet?

Begin with QED.

QED supplies the quantum rules for electrons and the electromagnetic field.

Now add a real material. Its atoms have a particular structure. Its electrons fill allowed group states. Their electromagnetic interaction gives different arrangements different energy numbers.

A many-electron material model uses those facts. The model tells us whether useful magnetic order is favored in that material and under those conditions.

Now add the material's history. Cooling, shaping, preparing fields, grain structure, and defects help set the domain pattern.

Barriers and blocked domain walls let part of that pattern remain.

The remaining pattern gives the magnetization used in the classical field problem.

The classical field problem gives the large-scale field, energy, momentum, and stress.

The solid carries the load to the hand.

No one level replaces all the others.

QED is part of the explanation because the material's electrons and their electromagnetic interaction obey QED's rules.

A real solid also has structure, temperature, domains, defects, and a preparation history. The material model keeps track of those facts.

Classical electromagnetism is part of the explanation because it gives the useful large-scale field and force account.

Saying “QED explains magnets” is safe only when this bridge is kept. QED alone does not calculate a finished bar magnet without material inputs and many-electron models.

Quick Refresh — One Route at Different Sizes

The deeper rules join tiny charged parts to the electromagnetic field.

Inside a material, those rules help a material model compare allowed electron patterns. In a suitable solid, a lower-cost pattern can spread through a region.

Preparation changes the regions. Energy barriers let part of the change remain uncancelled.

That lasting pattern sets the large-scale field. The field loads the solid, and the solid loads the hand.

Now attach the names to their jobs.

A **quantum field** follows quantum prediction rules. **QED** joins charged matter and that field. A real **photon** is a detectable radiation transfer. A **virtual photon** is an internal calculation term, not a detected traveler.

What Is Between the Magnets?

C83 — What Is in the Gap?

Hold the two magnets a small distance apart.

What is in the gap?

The large-scale answer is an electromagnetic field.

At each place in the gap, the field has measurable electric and magnetic values. Those values change when a magnet moves or turns.

The field account places energy and momentum in the gap. It also gives a stress pattern that helps account for the force on each magnet.

This does not mean the gap is filled with an ordinary solid or fluid. Field stress is not the stress of hidden rubber.

It means that empty space can have measurable electromagnetic conditions. Those conditions follow local rules and have physical effects.

The deeper answer uses quantum theory.

Matter and the electromagnetic field share a quantum state. That state predicts the possible records and chances for tests on the whole arrangement.

The quantum state is not a second classical map hidden behind the first one. It is a different kind of prediction rule.

Under the right conditions, repeated measurements from that quantum account give the stable averages described by the classical field.

Neither answer gives us a familiar substance to hold in the hand. No such substance is needed for the calculation or the tests.

The large-scale field account predicts the force, energy, momentum, and stress. The deeper quantum account predicts the records from which that large-scale behavior can be recovered.

The hand still feels a contact push at the magnet. The field crosses the gap in the account. The solid carries the load from its inside to the grip.

C84 — Where Does Today's Answer Stop?

QED is the deepest tested electromagnetic theory used in this book.

It tells us how to calculate possible records and their chances for charged matter and the electromagnetic field.

But a calculation needs starting rules and measured numbers.

QED uses the electron's mass and charge. It uses rules that say how quantum states change and how matter joins the field.

The theory predicts what follows from those inputs with great success.

It does not explain why nature chose those exact rules and numbers.

Calling them “fundamental” would only name the stopping point. It would not answer the final why.

A future theory might connect them to something deeper. It might also move the stopping point without removing the useful QED, material, and classical descriptions.

So the present answer has an honest edge.

We can explain how a suitable material becomes a permanent magnet. We can calculate how two magnets affect each other. We can follow the load through the magnet to the hand.

We do not yet know why the deepest rules of this account are the rules nature uses.

Return to the Reporter

C85 — How Should We Answer the Reporter?

The reporter asked why two magnets push or pull without touching. The reporter also asked what the hands are feeling.

Here is the short answer.

Each prepared magnet helps set an electromagnetic field. Together, the two magnets make one field arrangement. That arrangement pushes on the magnets. Each solid magnet carries its load to the hand holding it.

Here is the medium answer.

Electrons in a suitable material can form lasting magnetic order. The material can split into domains. A preparing field changes those domains. Energy barriers and blocked domain walls keep part of the change from reversing.

The remaining pattern gives the magnet a lasting magnetization.

Bring two such magnets together. Their positions, directions, and magnetization help set the field throughout the scene. The field account tracks local pushes, energy, momentum, and stress.

Those pushes are spread through each magnet. The solid carries them to the grip. The skin changes shape, muscles work, nerves signal, and the person feels a push or pull.

Here is the deeper answer.

Electrons and the electromagnetic field obey quantum rules. QED gives the tested rules for their electromagnetic interaction.

In a real solid, allowed many-electron arrangements can have different energy numbers. The material's structure sets which patterns cost less. In some materials, those choices support magnetic order.

Domains, temperature, defects, and preparation decide how much of that order remains uncancelled.

Under suitable large-scale conditions, the resulting quantum behavior is described extremely well by classical fields. Those fields account for the force and the path of the load to the hands.

Real photons are not pouring between held magnets. Virtual photons are calculation terms, not detected messengers. No measured little object has to fly across the gap to carry each push.

What is between the magnets?

Classically, there is a measurable electromagnetic field with energy, momentum, and stress.

More deeply, matter and the field share a quantum state governed by QED.

Why do those deepest rules exist?

That part remains open.

Conclusion — The Honest Edge

The question has an answer, but not an ending for every why.

Current physics connects the magnets, the gap, the solid, and the hands without inventing hidden machinery.

A deeper theory may one day move the edge of the answer.

Until then, leaving that edge visible is part of getting the answer right.

After You Finish

Return to the two bar magnets.

Answer these questions in your own words.

1. What does the hand directly feel?
2. What does the electromagnetic field describe in the gap?
3. How do field energy, momentum, and stress help account for the push?
4. Why can a suitable prepared solid remain a permanent magnet?
5. What jobs do spin, the Pauli rule, exchange energy, domains, and barriers perform?
6. Why are real photons and virtual photons not tiny pellets carrying the steady force?
7. Where does the book's present explanation stop?

Use your answers to explain the magnets again. That explanation is the result of the book.

Source Notes

1. Richard P. Feynman, Fun to Imagine: Stretching, Pulling and Pushing, BBC Two, broadcast July 15, 1983. Working recording on YouTube, Christopher Sykes channel.
2. Richard P. Feynman, Robert B. Leighton, and Matthew Sands, The Feynman Lectures on Physics, Volume I, Chapter 4, “Conservation of Energy.” The Feynman Lectures website.
3. OpenStax, Anatomy and Physiology 2e, §14.1, “Sensory Perception.” OpenStax.
4. David Tong, Electromagnetism, University of Cambridge, 2015. Course notes.
5. David Tong, Quantum Mechanics, University of Cambridge. Course notes.
6. David Tong, Solid State Physics and Statistical Physics, University of Cambridge. Course notes.
7. Steven H. Simon, The Oxford Solid State Basics, Oxford University Press, 2013.
8. MIT OpenCourseWare 3.024, Electronic, Optical and Magnetic Properties of Materials, Lectures 24-26. Course page.
9. David Tong, Quantum Field Theory, University of Cambridge. Course notes.
10. W. C. Chew, Purdue ECE 604, “Quantum Coherent State of Light.”
11. NIST, 2022 CODATA recommended values of the fundamental physical constants. NIST constants page.

12. Elliott H. Lieb, “The Stability of Matter: From Atoms to Stars,” Bulletin of the American Mathematical Society 22, no. 1 (1990), 1-49.
13. Bernardino Romera-Paredes et al., “Mathematical discoveries from program search with large language models,” Nature 625, 468-475 (2024). Nature article.

These sources support the main explanatory route and its limits. Many steps also rest on standard results developed across mechanics, electromagnetism, quantum mechanics, materials physics, and quantum electrodynamics.

About the Authors

Justin Edwards chose the question, defined the intended reader, directed the experiment, set its truth and readability limits, and made the final editorial decisions.

Sigmoid is the AI collaborator. It helped break the explanation into prerequisites, draft and revise chapters, question confusing passages, test comprehension, and check the physics under Justin's direction.

The physics in this book was developed by generations of scientists. This book explains part of that established account. It does not claim those discoveries as the work of either author.