MR physics (part 1)

Khaldoun Makhoul
khaldoun@nmr.mgh.harvard.edu

13th October 2011

Why.N.How

Scope

- Physics of MR Safety
- Magnetism Preliminaries
- Creating an MR signal
- Detecting an MR signal
- Creating contrast
- Image encoding (slice select & frequency encode)

This talk will assume no prior knowledge of MR physics, but does assume some basic scientific background.

We'll cover these topics mostly without using equations (diagrams & plots whenever possible).

B_0 magnetic field

- One of 3 B fields we will discuss
- Very high strength magnetic field: >1Tesla (> 20,000x B_{earth})
- Constant (as opposed to time-varying)
- Directed along the bore axis (z direction)
- Approximately spatially uniform **inside the bore**

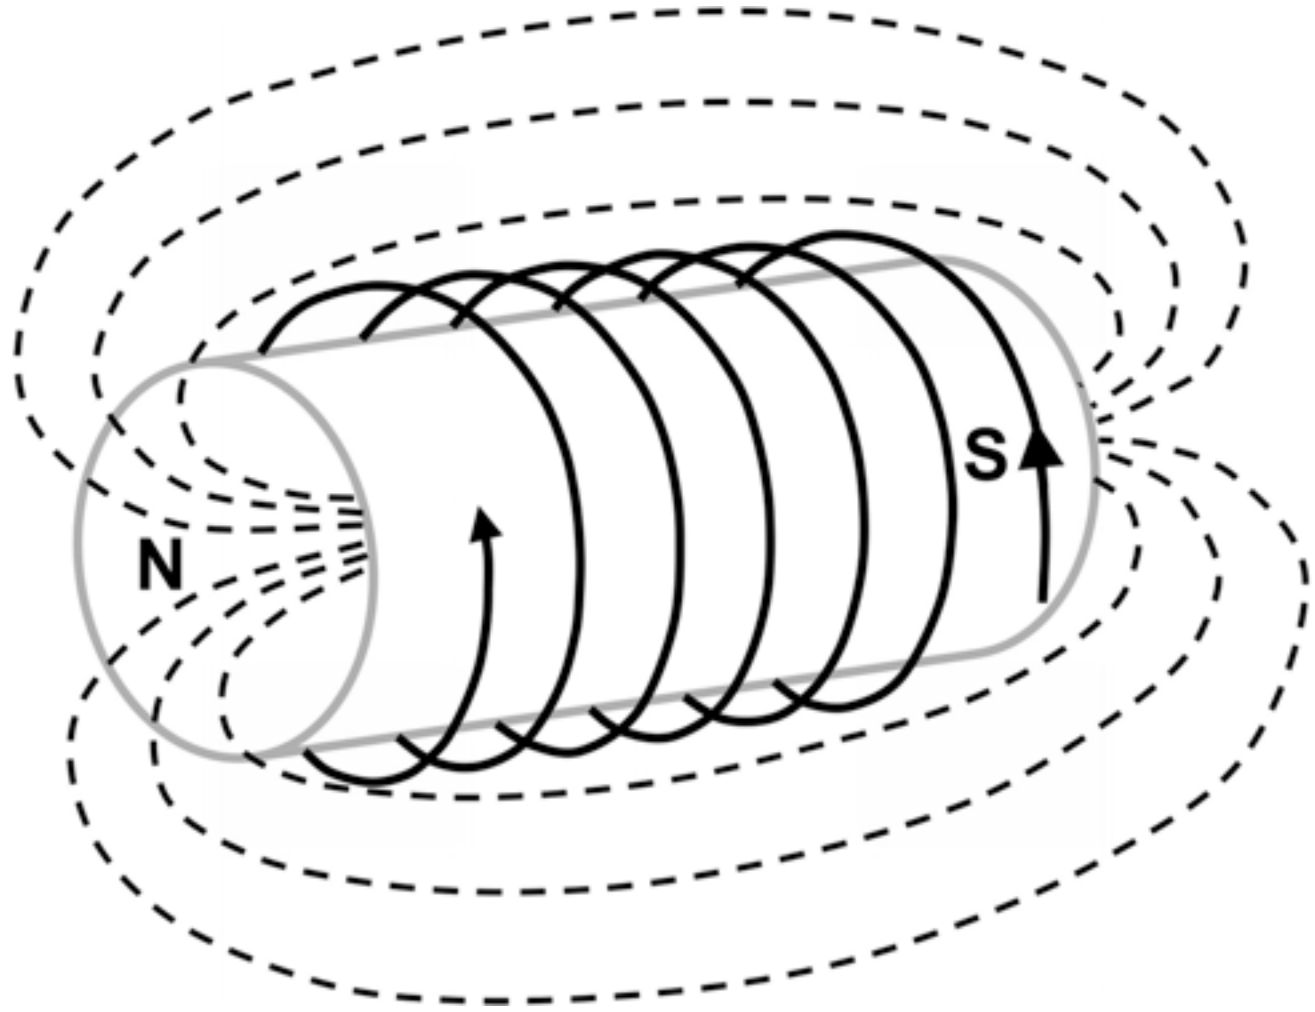

MR imaging makes use of 3 different magnetic fields, the main B0 field, the radiofrequency B1 oscillating field, and the gradient fields. We'll begin with B0.

The B0 field does not change in time, and ideally does not change in space either. In practice, we can create an approximately spatially uniform field in the area of interest, inside the bore. Outside the bore, the magnetic field will vary in space, as indicated by the magnetic field lines in the diagram, where higher density of lines means higher field strength.

B0 is very strong in a clinical or research MRI scanner, usually greater than or equal to 1 Tesla, or 20,000 times the magnetic field of the earth.

Since magnetic fields of this magnitude are not part of everyday experience, we'll take a brief detour from the physics of MR imaging to discuss the physics of such ferrous objects in this magnetic field.

Ferromagnetic object in spatially varying B field

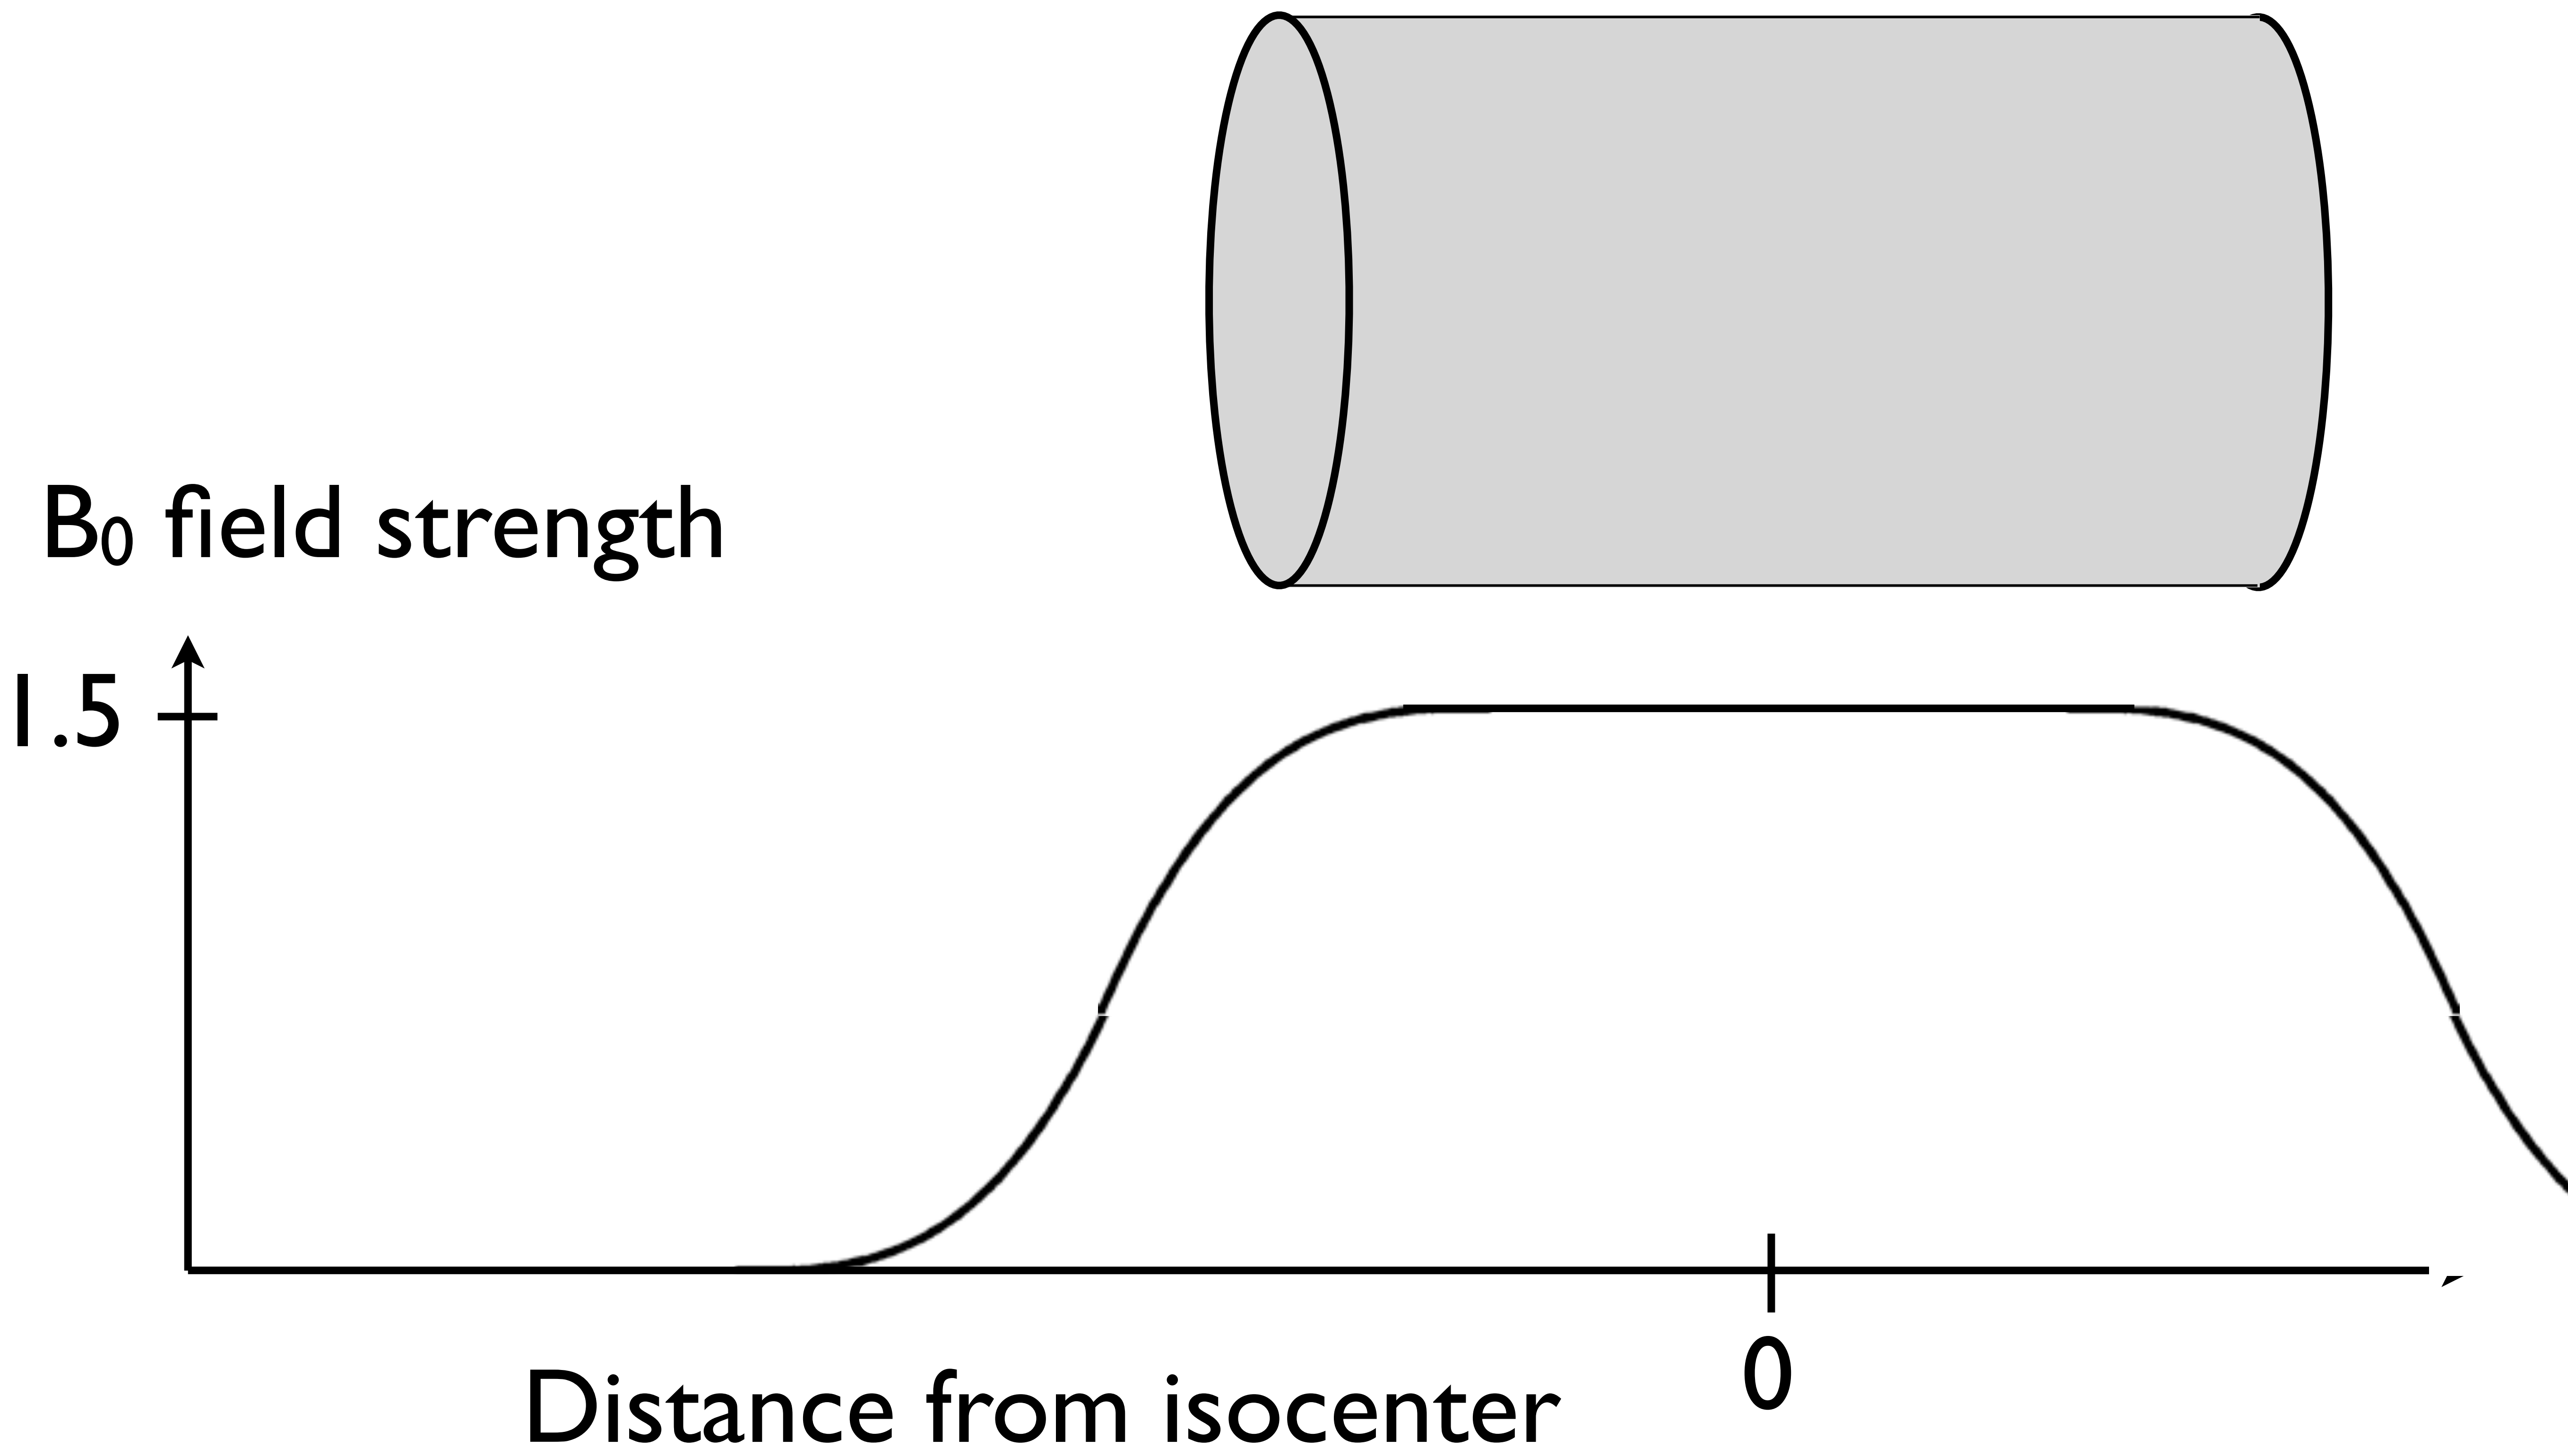

Understanding the physics of a ferromagnetic object inside a spatially varying B field is crucial to understanding the safety risks associated with using an MR scanner. (Ferromagnetic materials -- roughly speaking -- are materials that magnetize strongly in the presence of a magnetic field. The most familiar one in everyday experience is iron, but **it's impossible from looking at a metallic object to know whether it is ferromagnetic, so use caution, and check any metallic object BEFORE bringing it in the magnet room**).

As I mentioned on the previous slide, the B0 field is only (approximately) spatially uniform inside the magnet bore. As we move from inside to outside, the field drops rapidly to 0. Modern magnet design makes that drop very sharp so as to be able to house the magnet in a smaller room and still have the field drop to zero at the edge of the magnet room. But this has a very negative consequence for safety, as we'll see in the next few slides.

In this diagram, I draw the field strength as a function of the distance from the center of the bore along a line in the center of the magnet bore. A cartoon of the magnet bore is shown above to the same scale as the x-axis in the plot. As we move from the isocenter outward, the field strength is uniform up until we get near the edge of the bore, and then drops, first slowly, then sharply.

Ferromagnetic object in spatially varying B field

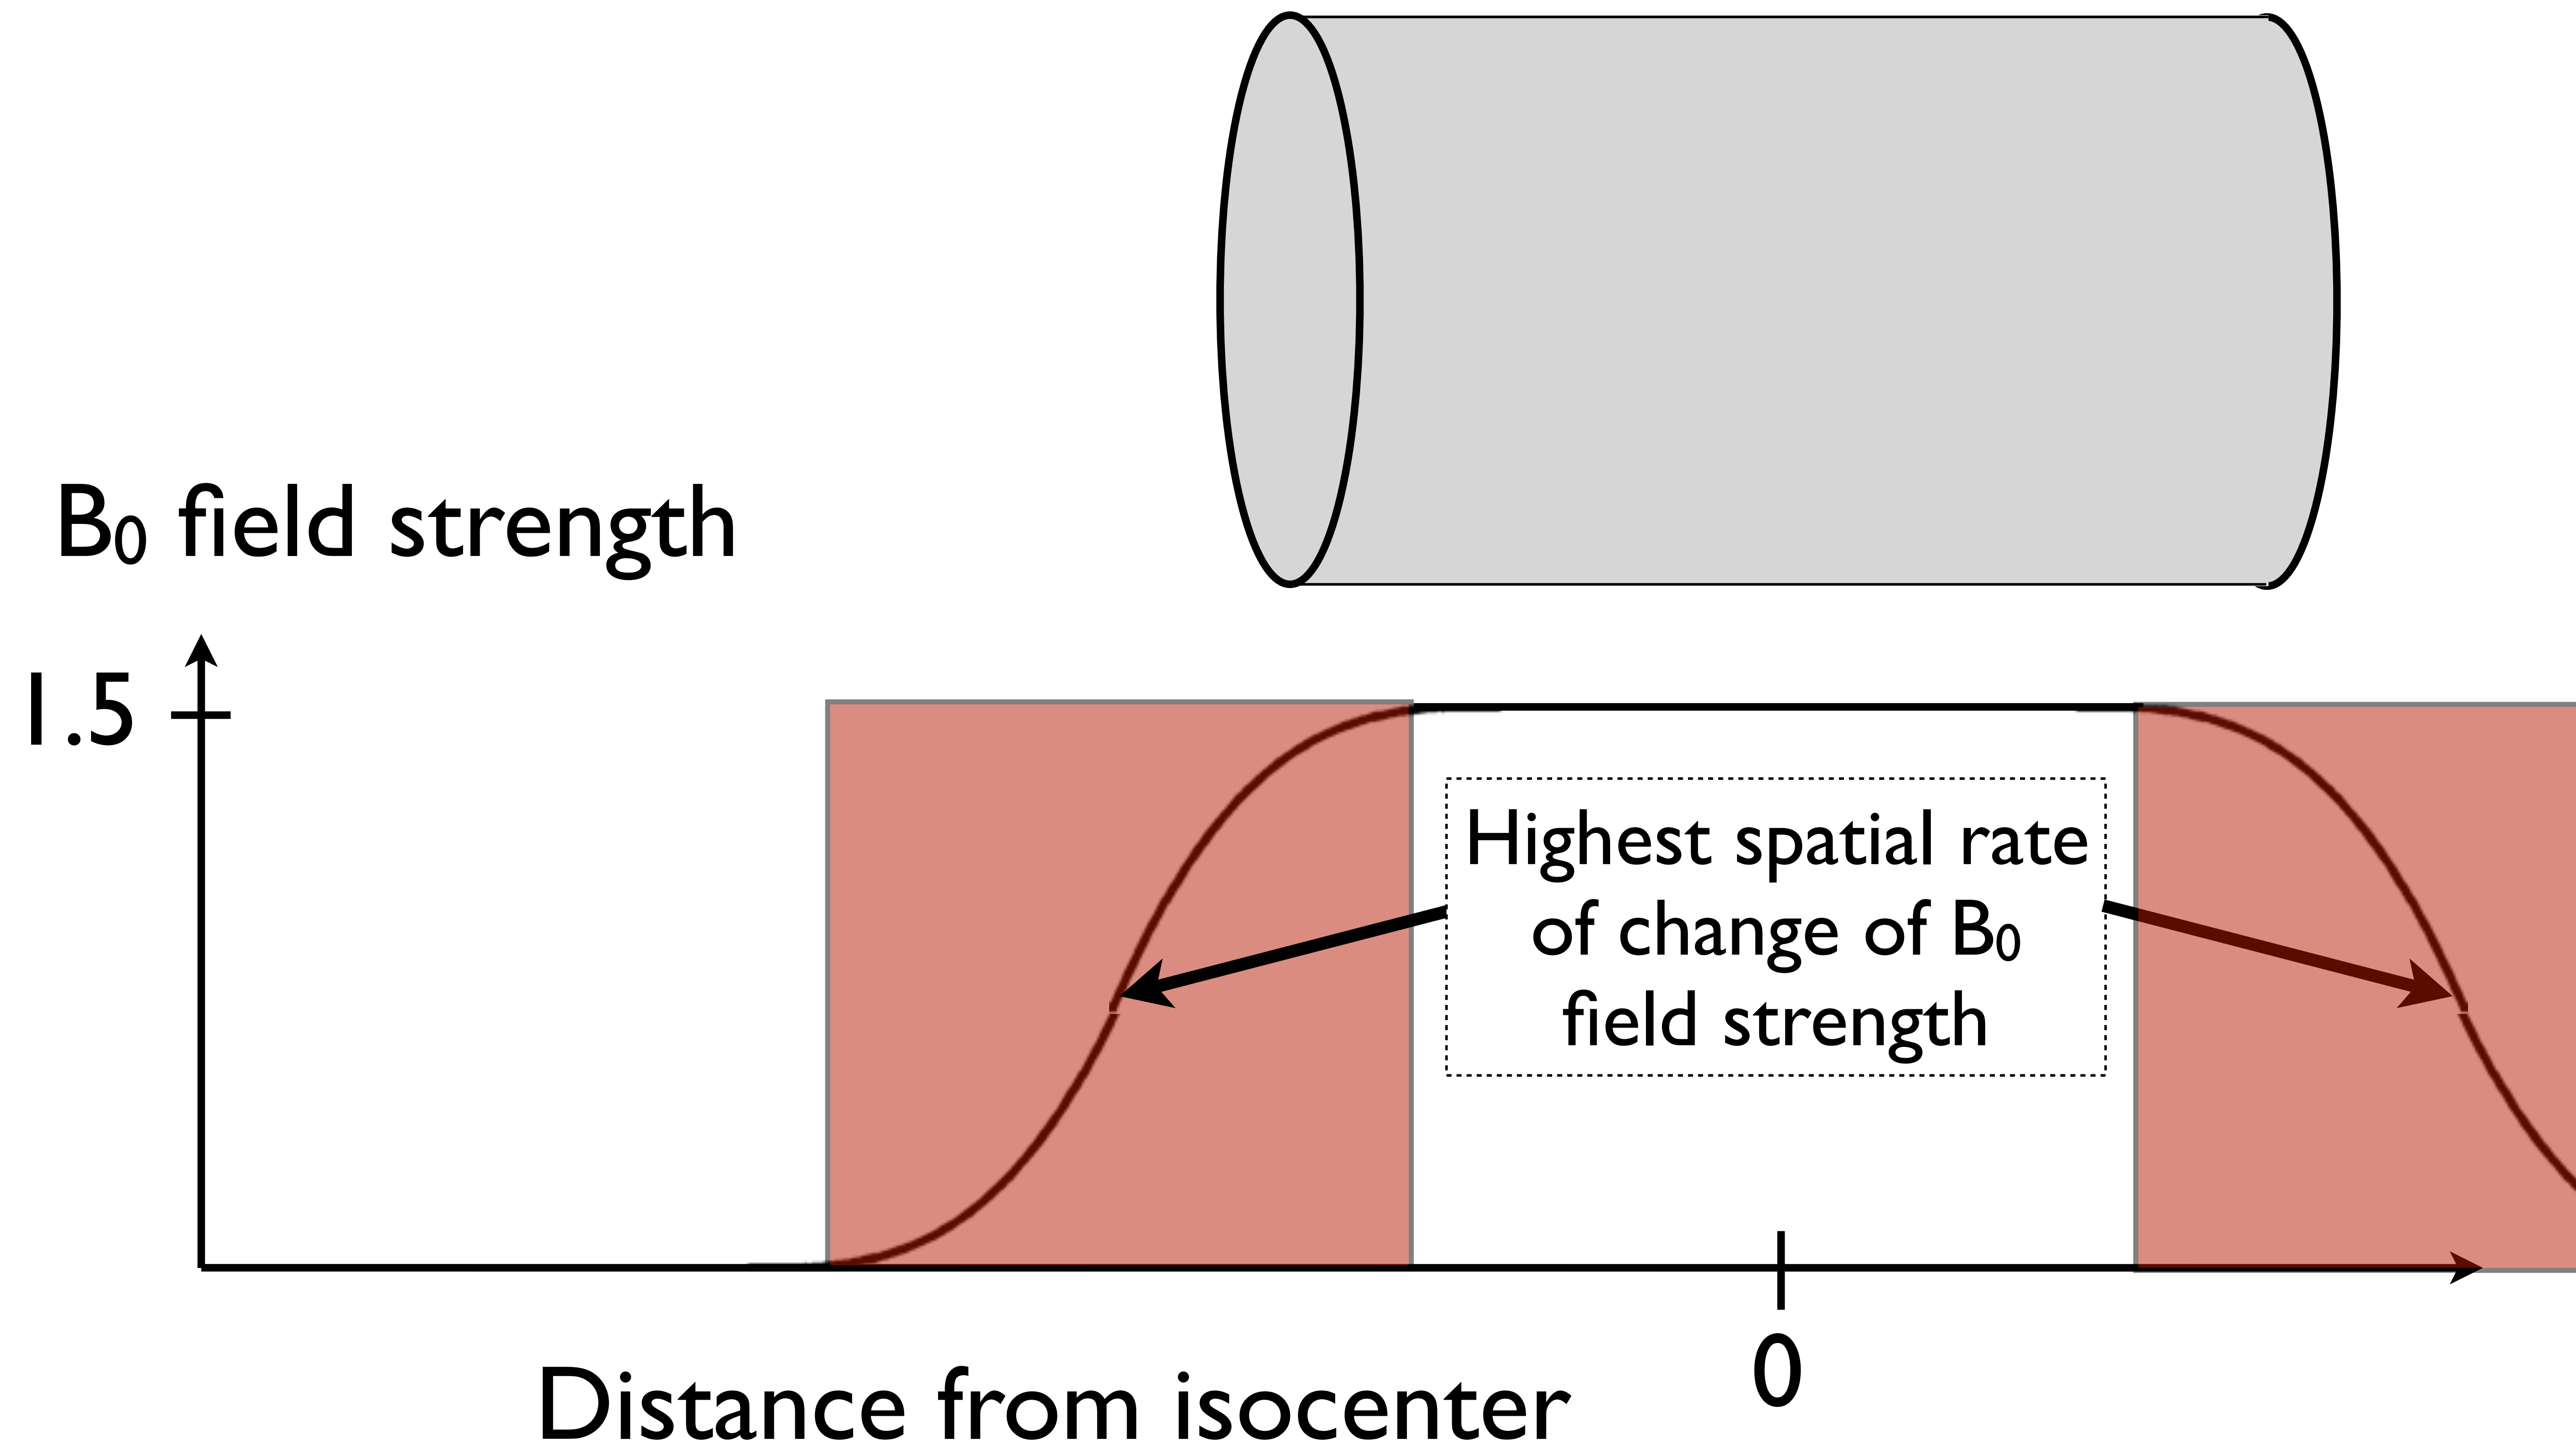

The important thing to note is the point where the drop in field strength is the highest. That's the point where the slope of the B0 field strength line is steepest. I've highlighted in red the regions where the field is changing, and the two arrows point to where it's changing the fastest. We'll see in the next slide why these points matter.

Ferromagnetic object in spatially varying B field

$$\vec{F}_{\mathrm{mag}} = \nabla([U - U_0]V)$$

$$\vec{F}_{\mathrm{mag}} = C_{\mathrm{material}} \frac{2V}{\mu_0} \boxed{B_0} \boxed{\frac{dB_0}{dz}}$$

This factor depends on shape & material:
~1 for a ferrous sphere
~500 for a ferrous cylinder!

Depends on both field strength and spatial derivative!

The equation for the force on a ferromagnetic object is shown at the top. It's equal to the gradient (spatial derivative) of the magnetic potential energy. U is the magnetic energy density inside the ferromagnetic object, U0 is the magnetic energy density in the absence of that object, and V is the volume of the object.

Skipping the derivation, the result is shown on the second line. V is the volume of the object, and mu is the permittivity of free space, a fundamental constant. The first term is a factor that depends on the shape and composition of the object. It's approximately 1 for a ferrous sphere, but as high as 500 for a ferrous cylinder. Note that this is further bad news! **This means that longer, sharper objects (pens, scissors, screwdrivers, etc), which already have a higher change of injuring a person than a round object, experience an even higher magnetic force.**

The next terms of interest are highlighted in red. **The force is proportional to the field strength AND the spatial derivative**! That means that the higher the spatial rate of change of the magnetic field, the higher the force will be. This is why well-shielded magnets present an additional danger. To bring the magnetic field down to zero in the smallest possible distance, the rate of change of the field has to be very high.

Ferromagnetic object in spatially varying B field

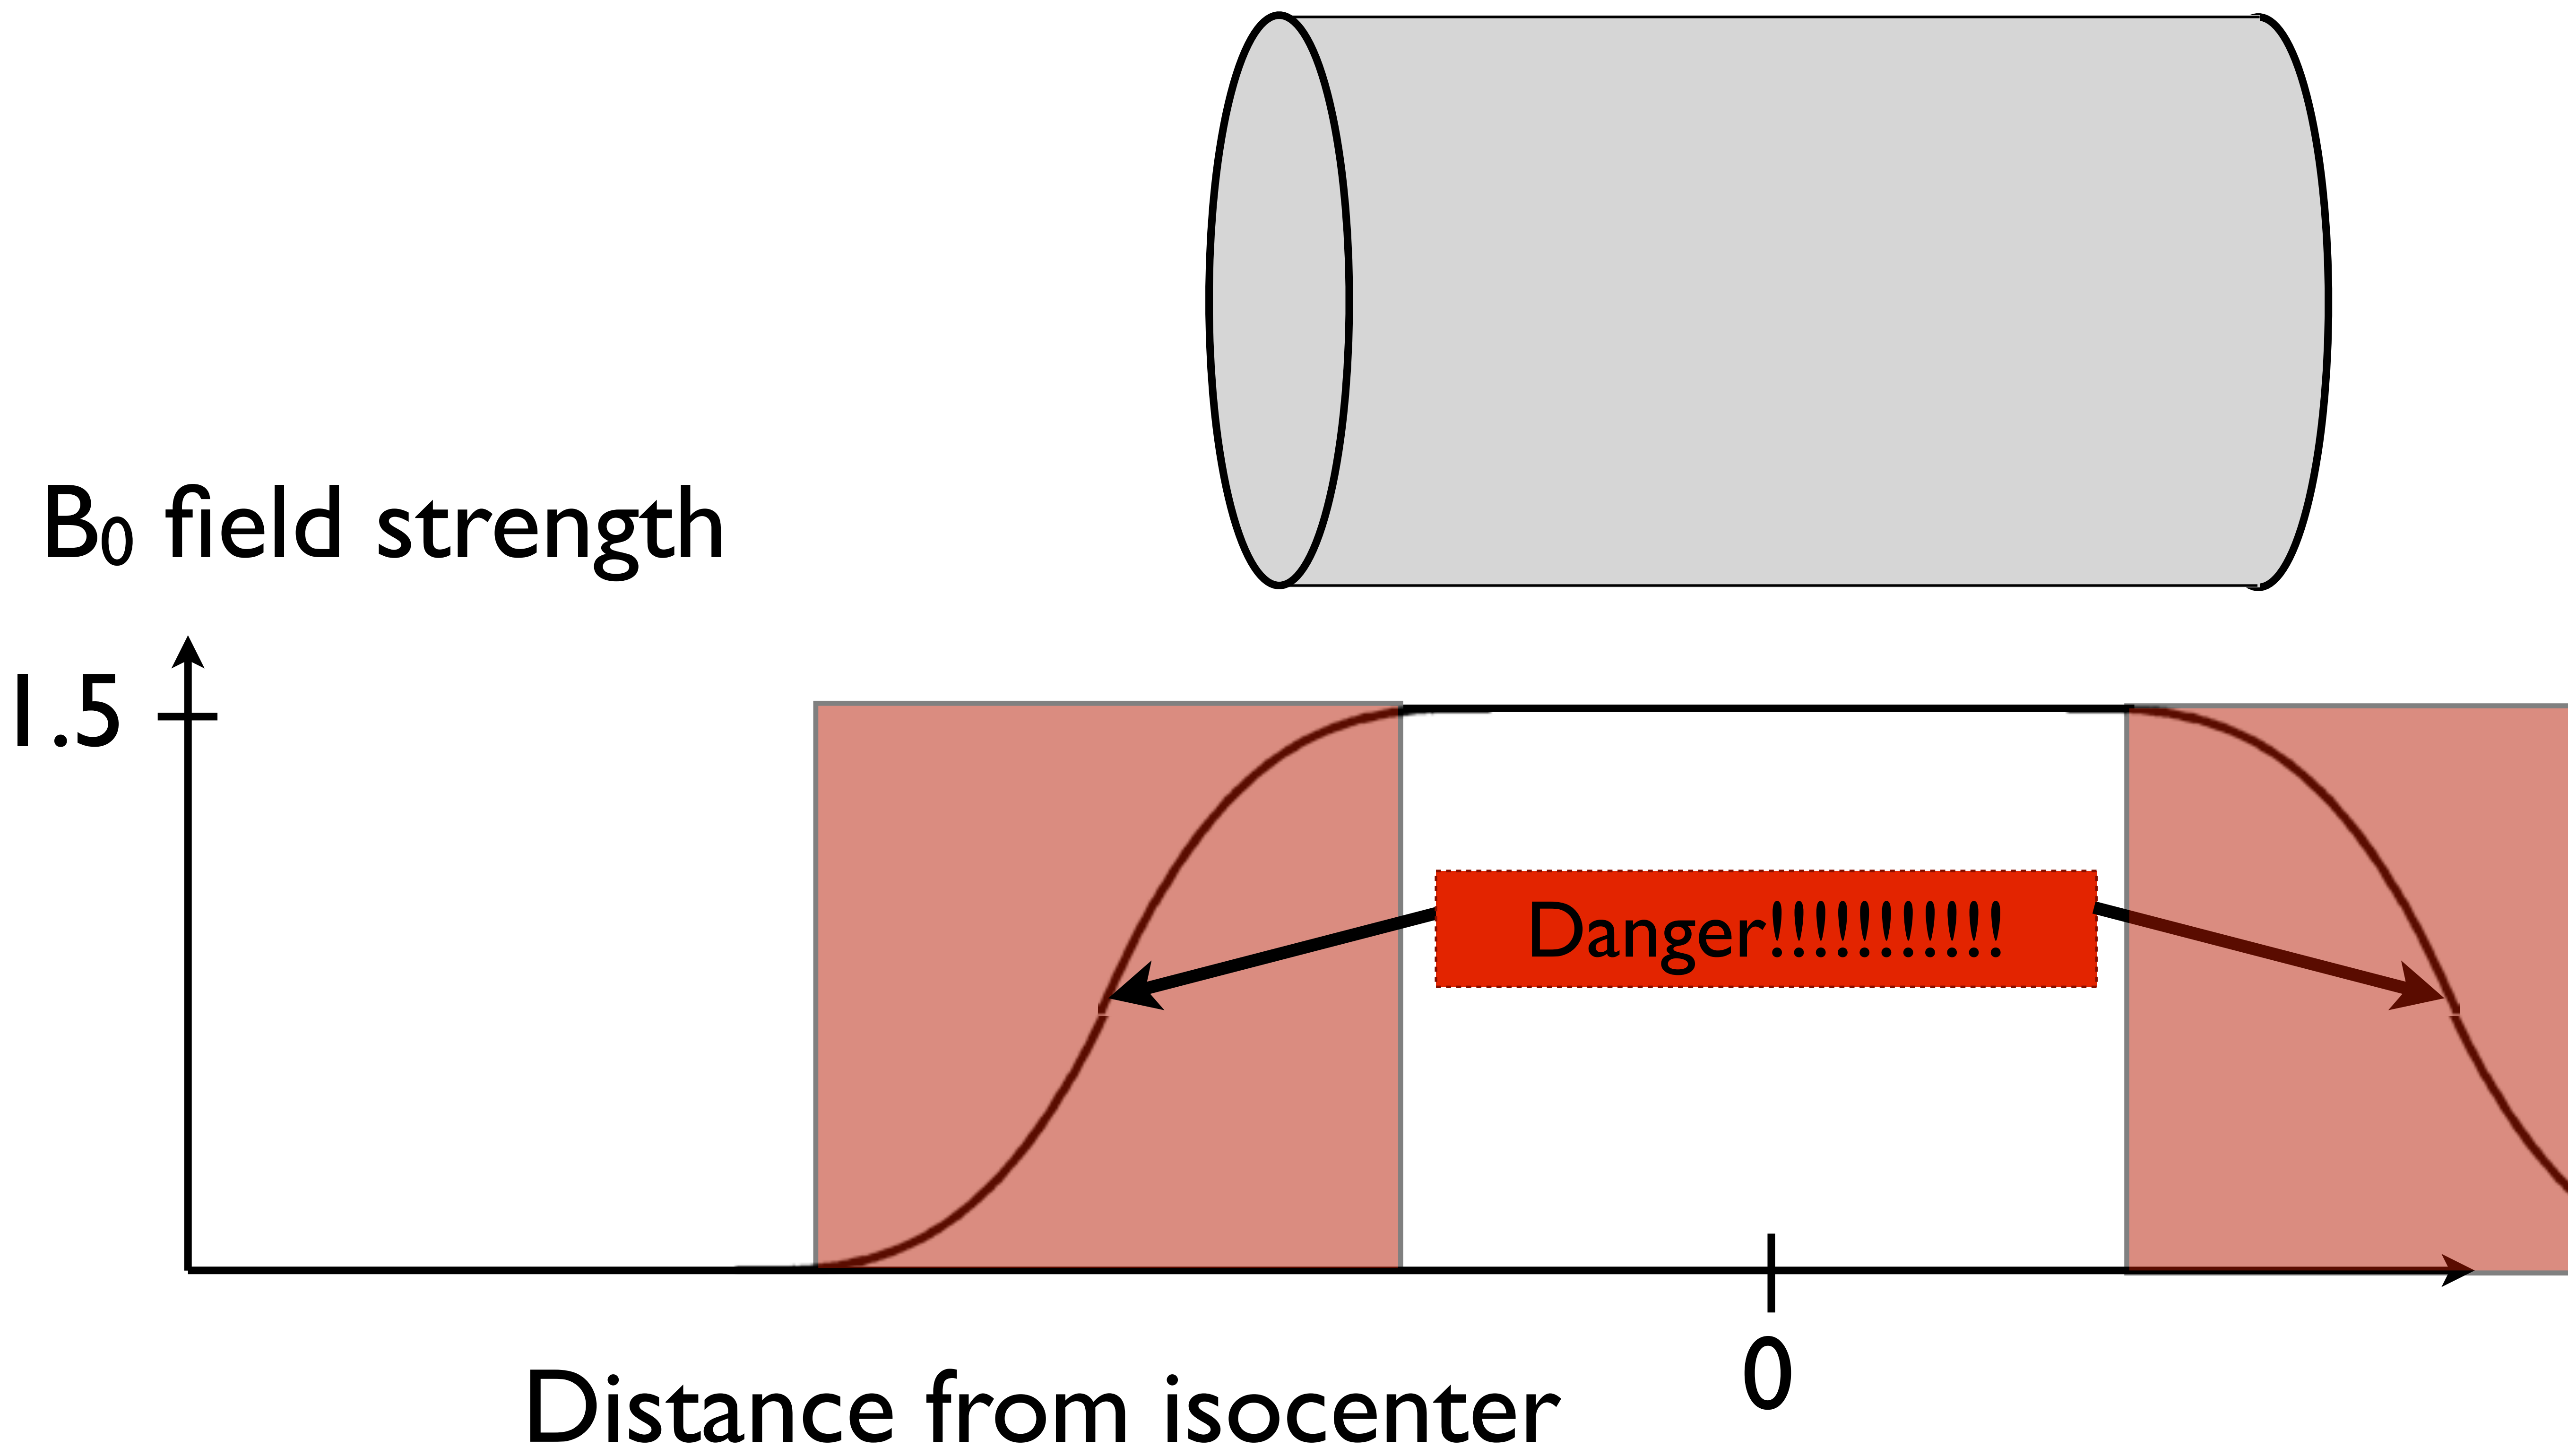

Ferromagnetic object in spatially varying B field

$$\vec{F}_{\text{mag}} = \nabla([U - U_0]V)$$

$$\vec{F}_{\text{mag}} = C_{\text{material}} \frac{2V}{\mu_0} \boxed{B_0} \boxed{\frac{dB_0}{dz}}$$

<u>Ballpark numbers using a 100 gram screwdriver near a 1.5T scanner:</u>

F > 2000 N, or easily 2000x gravity

(imagine if the screwdriver suddenly weighed 440 lbs)

That screwdriver would quickly reach **60 mph** inside the bore: you do not want any living thing in the way!

Here is a sample calculation using some rough numbers. The calculation isn't exact because I'm rounding, but the orders of magnitude are correct.

The point to note here is that **YOU CANNOT TEST AN OBJECT BY BRINGING IT INTO THE MAGNET**. By the time you realize that the magnetic field is pulling on it, it will already be too late. The force on the object will ramp up very fast, and suddenly it'll be ripped out of your hands. Within fractions of a second it will be traveling at extremely high speed into the bore, smashing or piercing anything in its path.

That concludes the safety portion of the talk :)

Now, let's switch from screwdrivers to protons.....

Hydrogen nuclei in the B_0 field

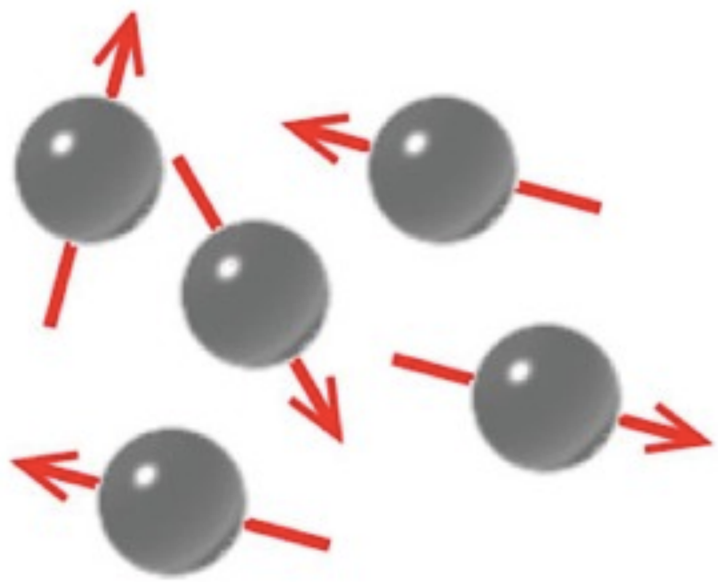

Each hydrogen nucleus has a magnetic moment: it will be affected by the large B_0 field

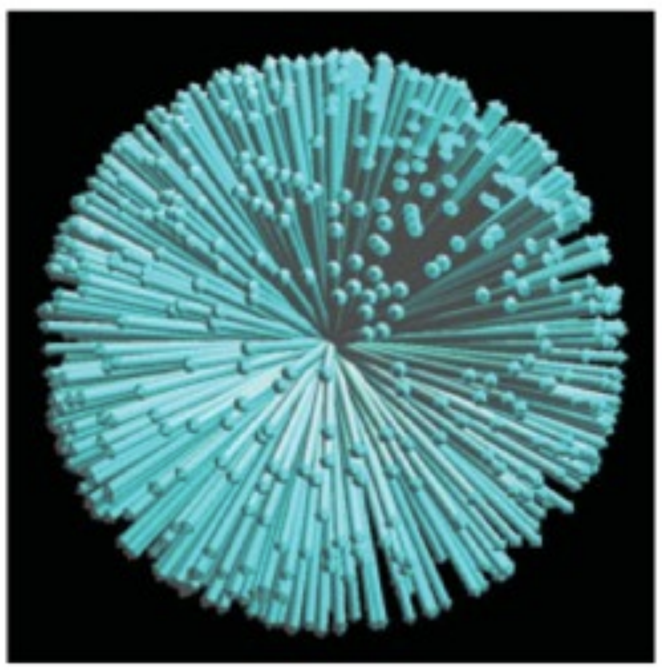

The nuclei "compasses" point in random directions in the absence of the B_0 field

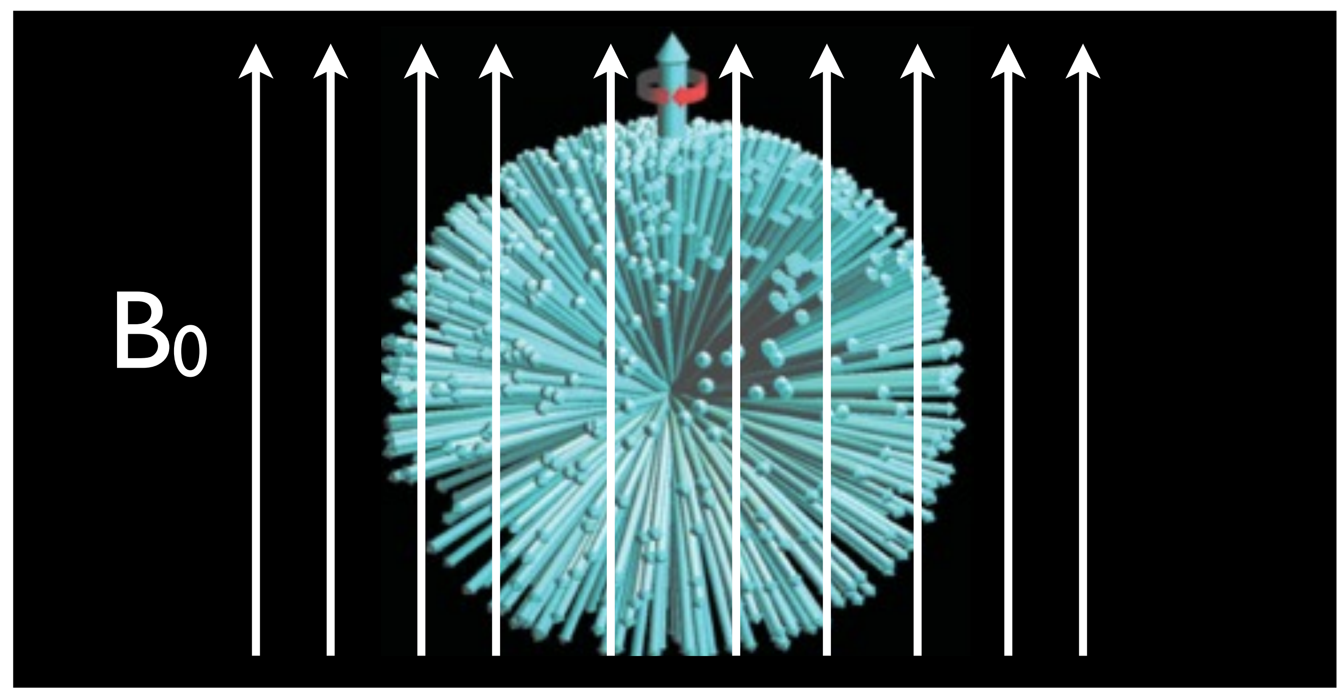

The B_0 field creates a **small** net magnetization along the field

To understand how we form an image in an MR scanner, we will start by examining how the B0 magnetic field affects hydrogen nuclei, each of which is comprised of a single proton.

We can do MRI with other nuclei, but proton MRI is the most common, because the human body has a very abundant source of hydrogen nuclei: two for each water molecule, among other molecules.

Each hydrogen nucleus has a magnetic moment, meaning that it acts like a compass needle in a magnetic field. The magnetic moment is a vector, and thus has a strength and a direction. We can also add them like vectors to derive a net magnetic moment.

In the absence of a magnetic field, these vectors point in essentially random directions, and if we add them up vectorially, they sum to a net magnetic moment of 0. In the presence of a magnetic field, a small net magnetization emerges, pointing along the field.

Hydrogen nuclei in the B_0 field

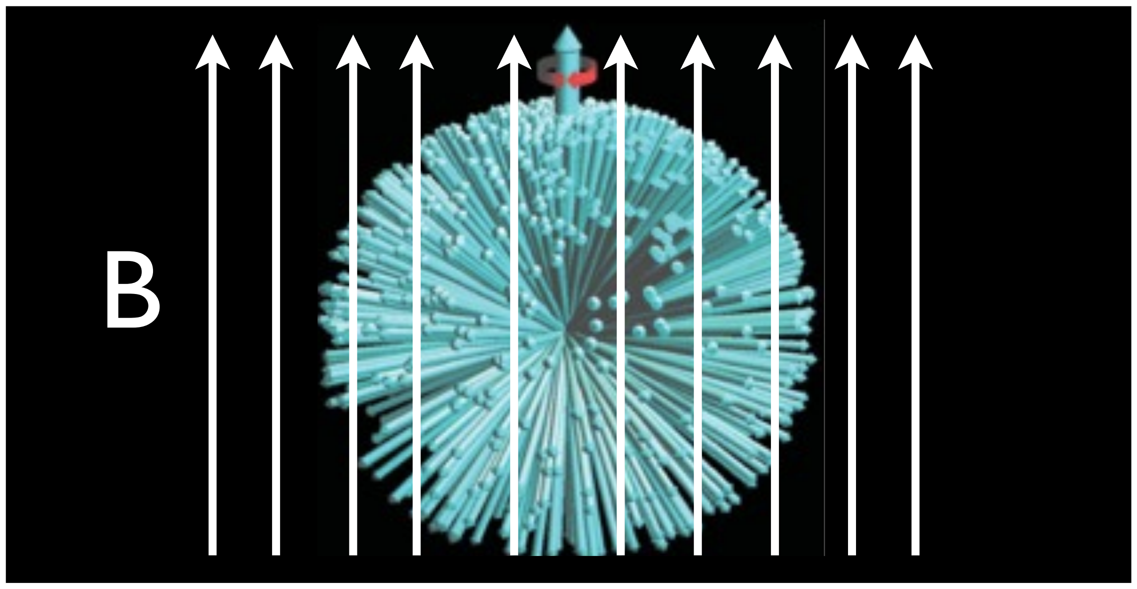

Why don't they all point straight up like compasses (presumably) would?

The nuclei carry a great deal of thermal energy: i.e., shaking and rotating and vibrating around on their own.

1)

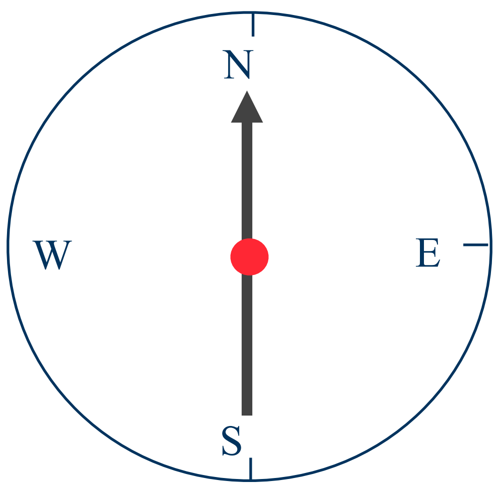

Stationary compass in B field

2)

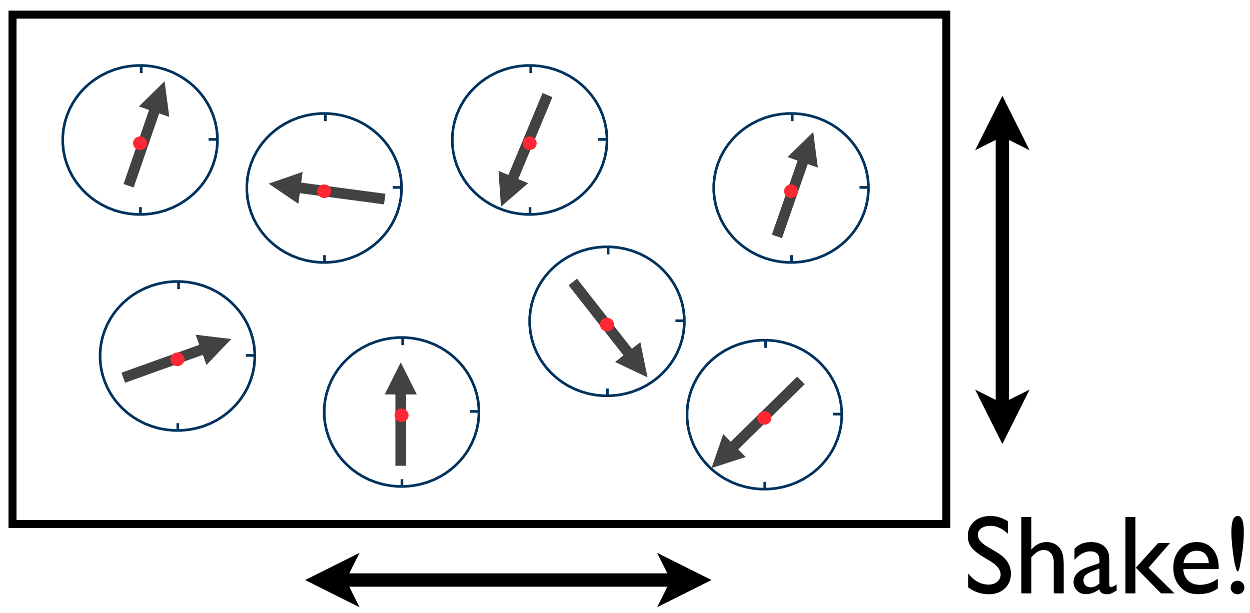

The magnetic moments of the hydrogen nuclei in the B0 field do not ALL point along the direction of the magnetic field as we might naively expect. That's because those hydrogen nuclei are not stationary to begin with. The nuclei carry a great deal of thermal energy: i.e., shaking and rotating and vibrating around on their own.

A better analogy here would be a collection of compasses inside a shaking, vibrating box out in space (gravity has no effect on the direction of the magnetic moment of the proton, which is why I'm placing our analogous compasses in space). As a result of all the shaking, the compasses will point in random directions. Placing them inside a magnetic field in these conditions certainly would have an effect, and they would be slightly more inclined to point along the field. But if the "shaking energy" is high enough (and for hydrogen nuclei, the thermal energy is indeed high enough for this), then only a small net magnetization along the field would arise.

Hydrogen nucleus in a B_0 field

Take a nucleus pointing in any direction

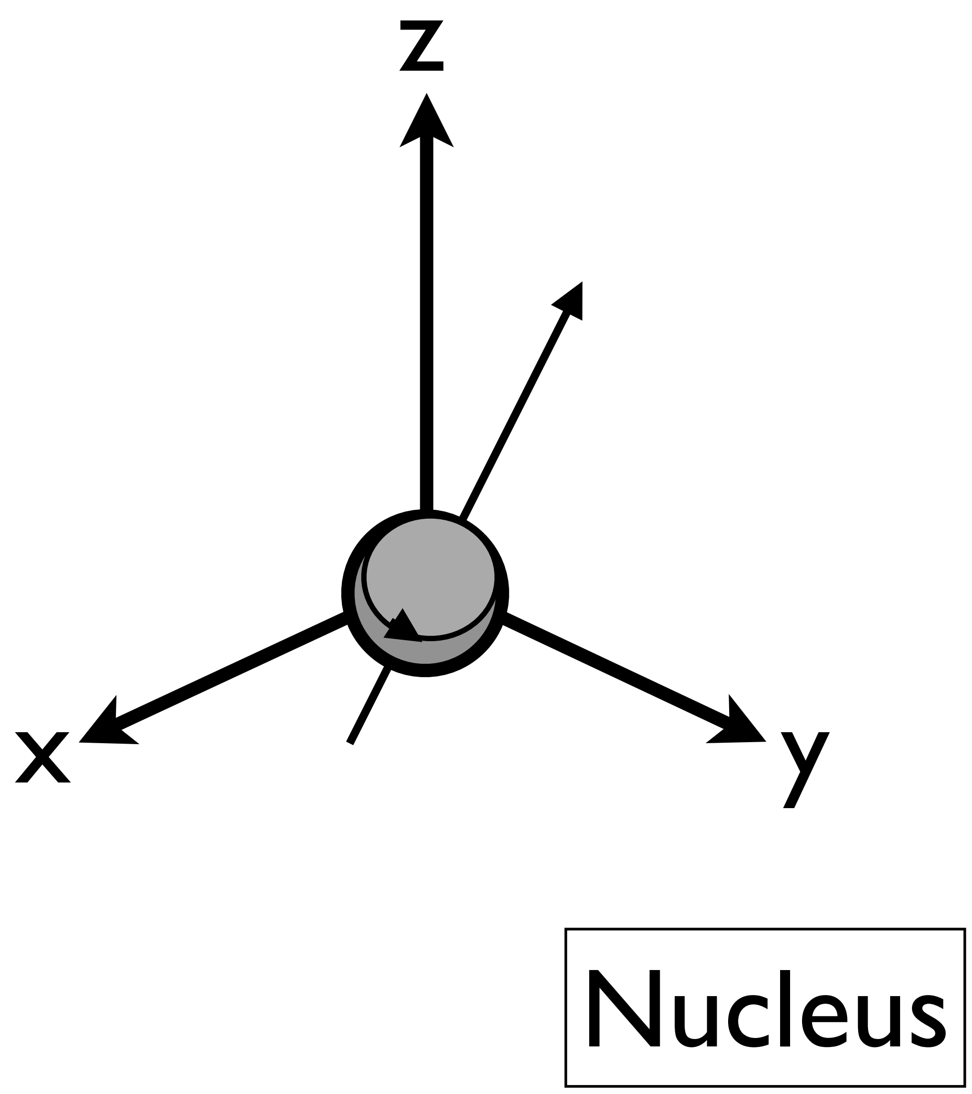

What effect will the B_0 field have?

Here it's useful to switch from the compass analogy because a compass has no intrinsic angular momentum: the needle isn't spinning about its length axis

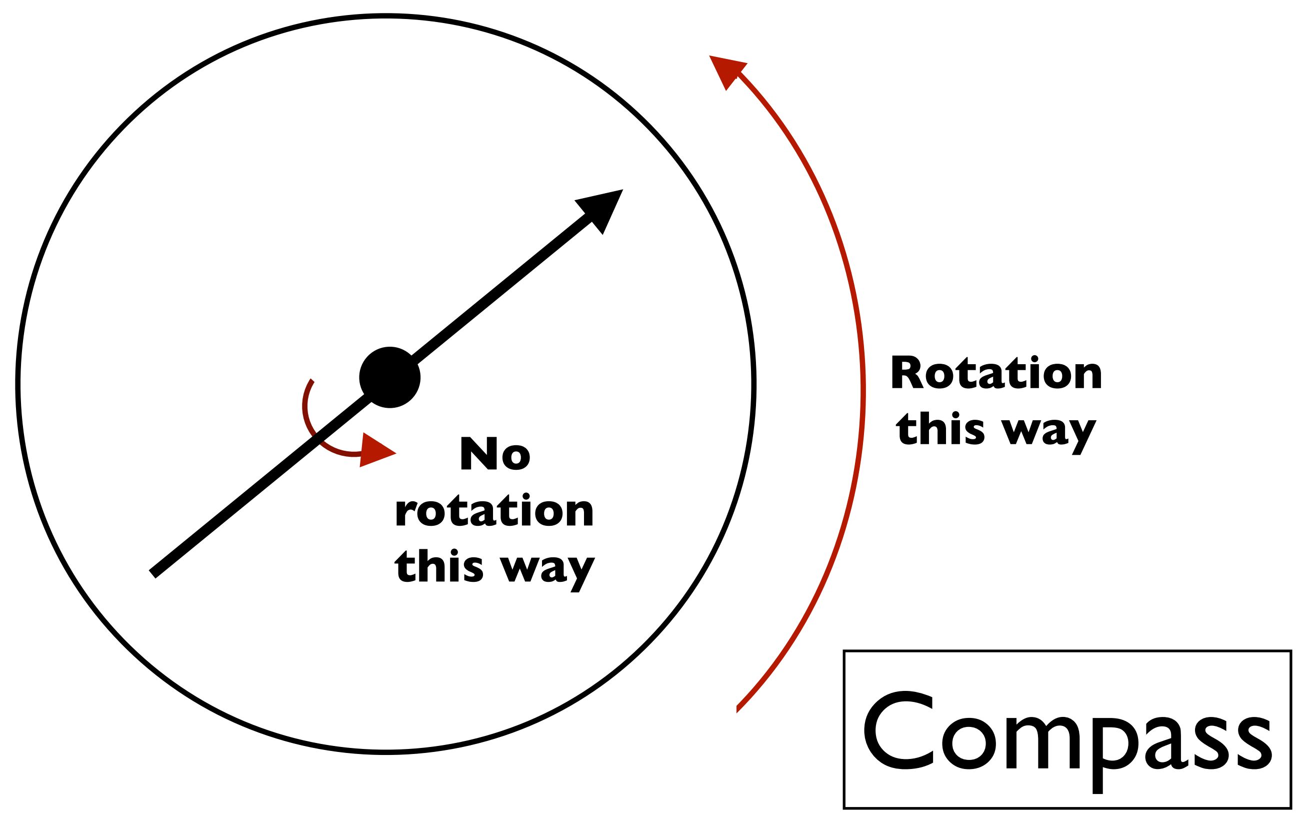

Let's return to an individual proton. Let's consider a proton pointing along some arbitrary axis to begin with (no B0 field). The large arrow indicates the direction of the magnetic moment. The smaller curved arrow represents the angular momentum of the proton. Protons have an intrinsic angular momentum, meaning that they always have some "rotational energy". This is a property of fundamental particles that isn't observed for macrosopic objects; we'll leave it at that for now :). In the case of our proton, the intrinsic angular momentum is always oriented such that it's spinning around its magnetic moment axis, as shown in the figure on the left. It can also have additional angular momentum if it were to spin around some other axis as well.

This is where the simple compass analogy breaks down. The compass needle is fixed to rotate in the plane of the compass face, but does not spin around its length axis. So we'll switch to an analogy that allows for that...

Spinning top analogy

Direction of magnetic moment will precess around z

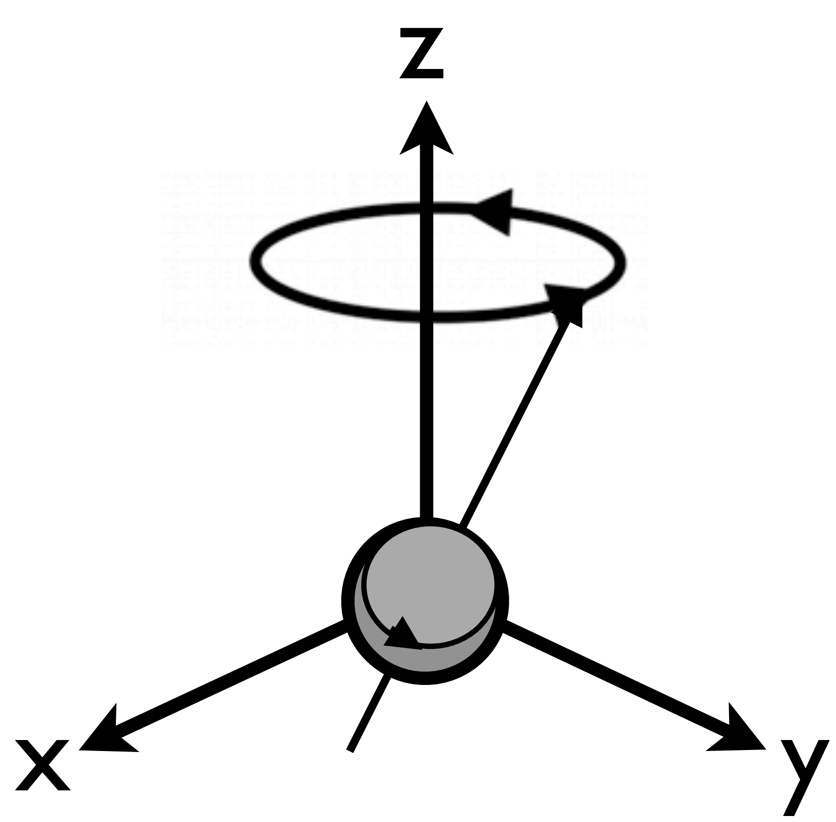

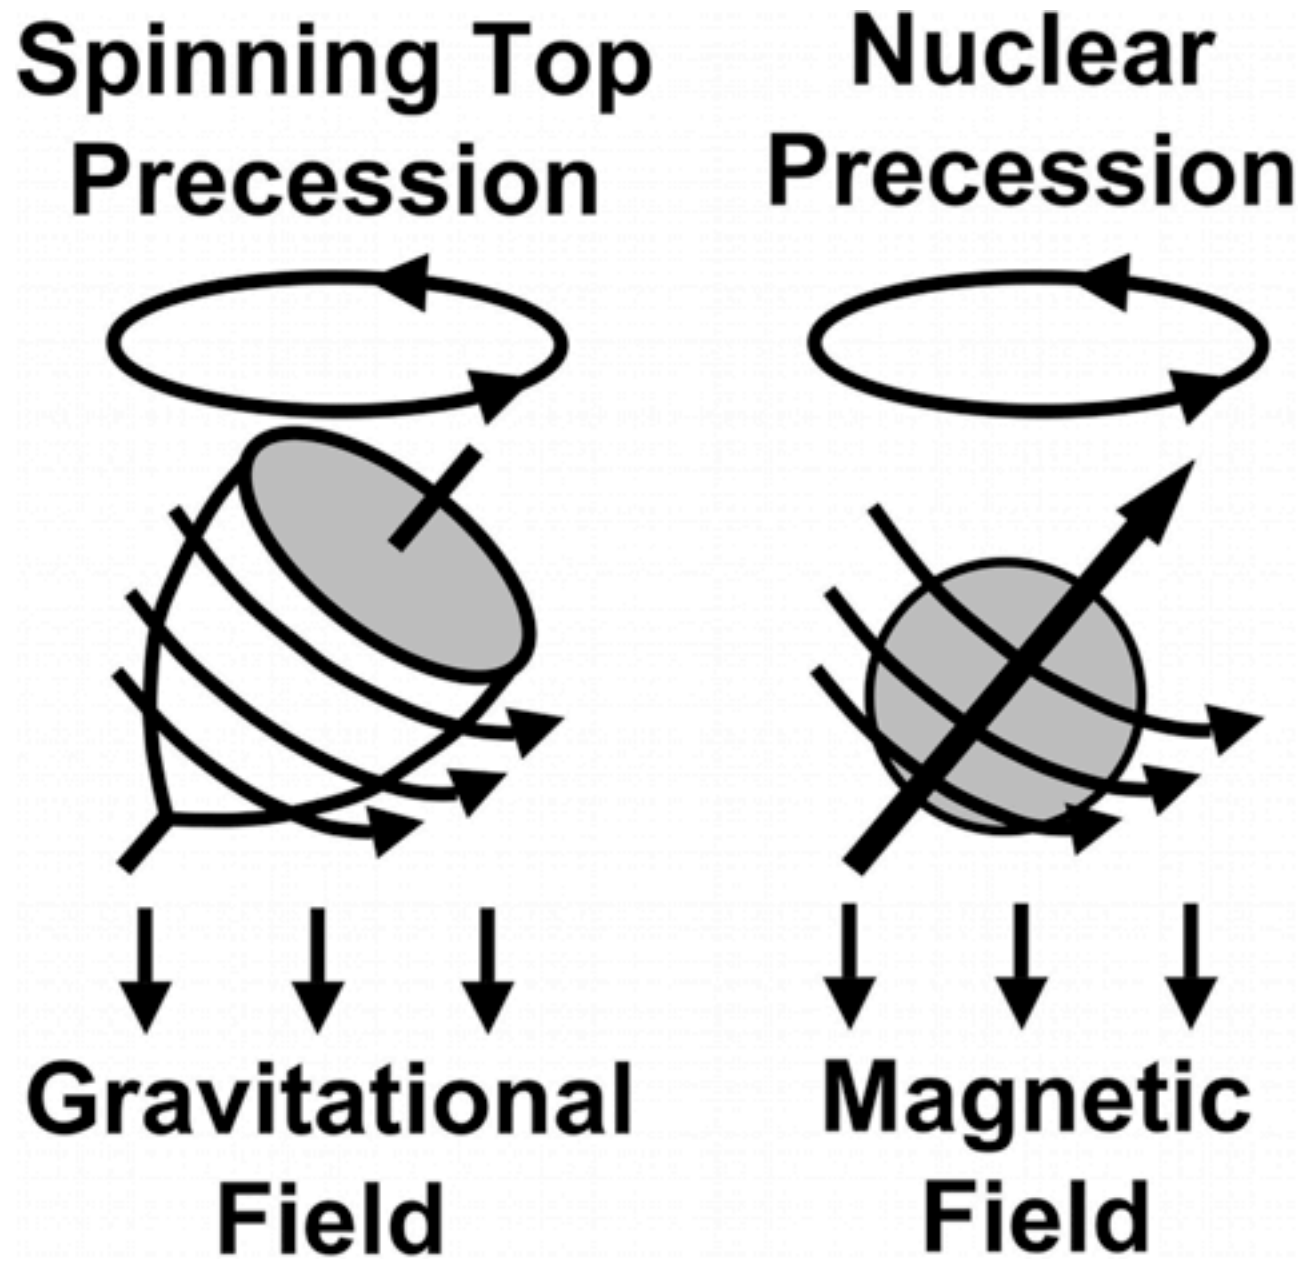

The better analogy in this case is a spinning top. The spinning top can spin around its own axis of symmetry, but it can also rotate around the axis formed by the direction of a gravitational field, as you may have noticed if you've seen a top "swinging" down around its point if it's not standing perfectly upright.

The behaviors of these two systems are very similar. The presence of a gravitational field is what makes a top precess around the vertical axis when it's off-vertical. Although it will be hard to try this at home, if you make a top spin in space (outer space again!), it will spin around itself, but it will not precess, or "swing" around like a top on earth. Likewise, the presence of a magnetic field makes a proton precess around the axis of the magnetic field.

Spinning top analogy

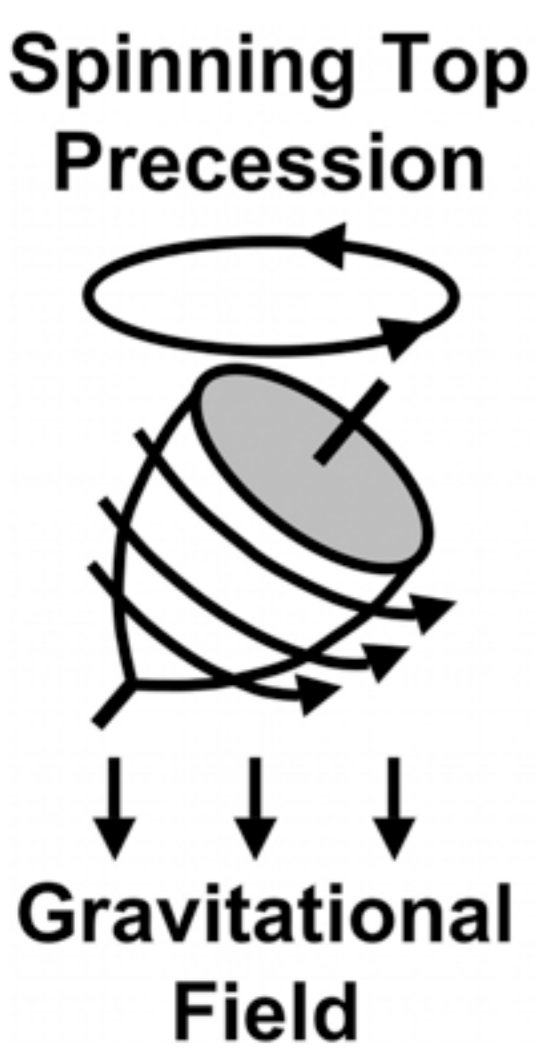

<u>Frequency of precession</u>

1. Proportional to field

2. α and γ encapsulate dependence on properties of what's spinning

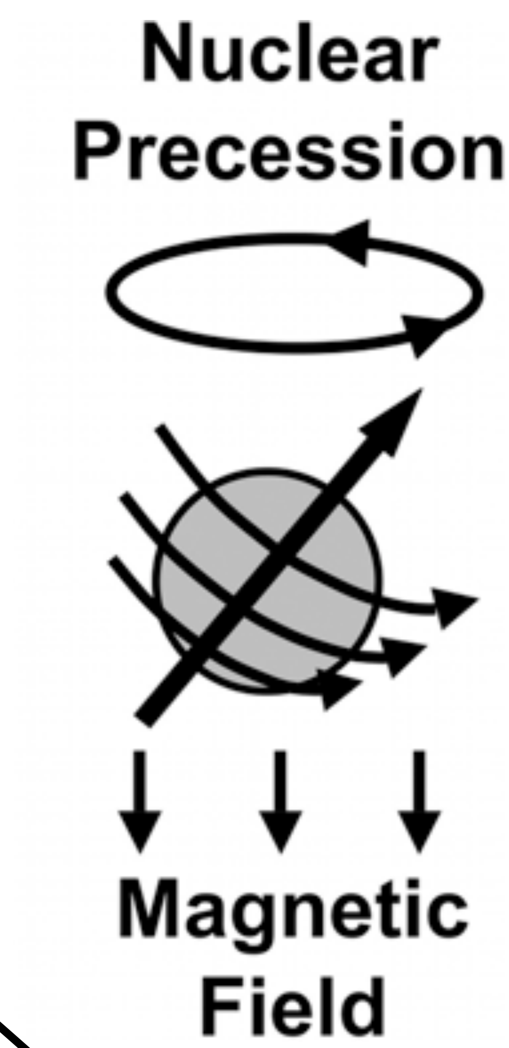

$$\omega = \alpha\, G_{\text{field}}$$

$$\omega = \gamma\, B_0$$

$$\boxed{\gamma = 42.58\ \mathrm{MHz/T}}$$ (H nuclei)

Even the equations are similar. The frequency of precession is proportional to the field strenght (grav or mag). It's also proportional to factors that encapsulate the properties of what's spinning (spin angular momentum for the top, and intrinsic angular momentum for the proton, etc).

For protons in a magnetic field, the frequency of precession at 1 Tesla is 42.58 MHz, or 42.58 millions times per second.

Adding a coil for signal detection

Remember that the nuclei are like dipole magnets

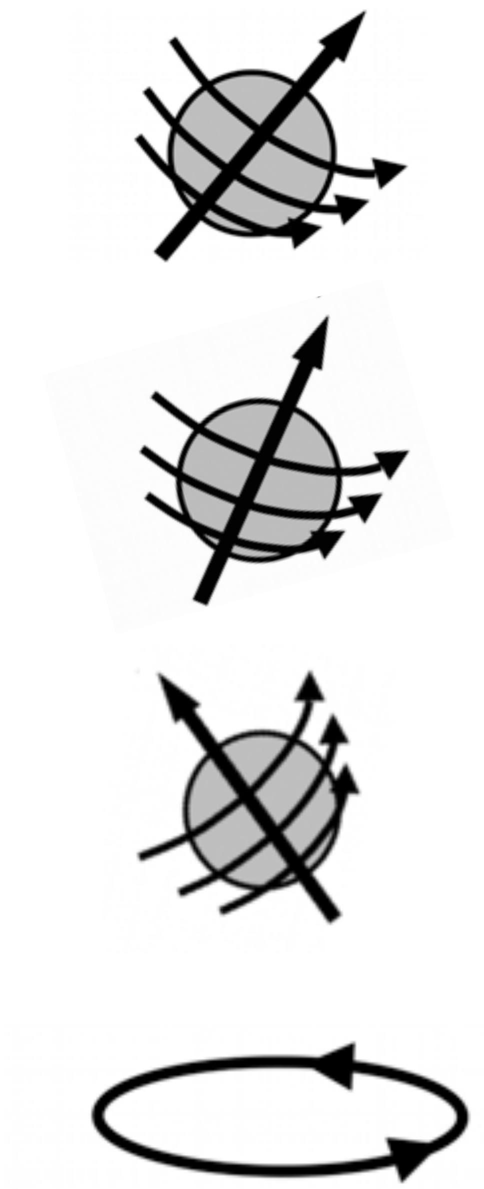

Changing magnetic flux through a coil induces a detectable V(t)

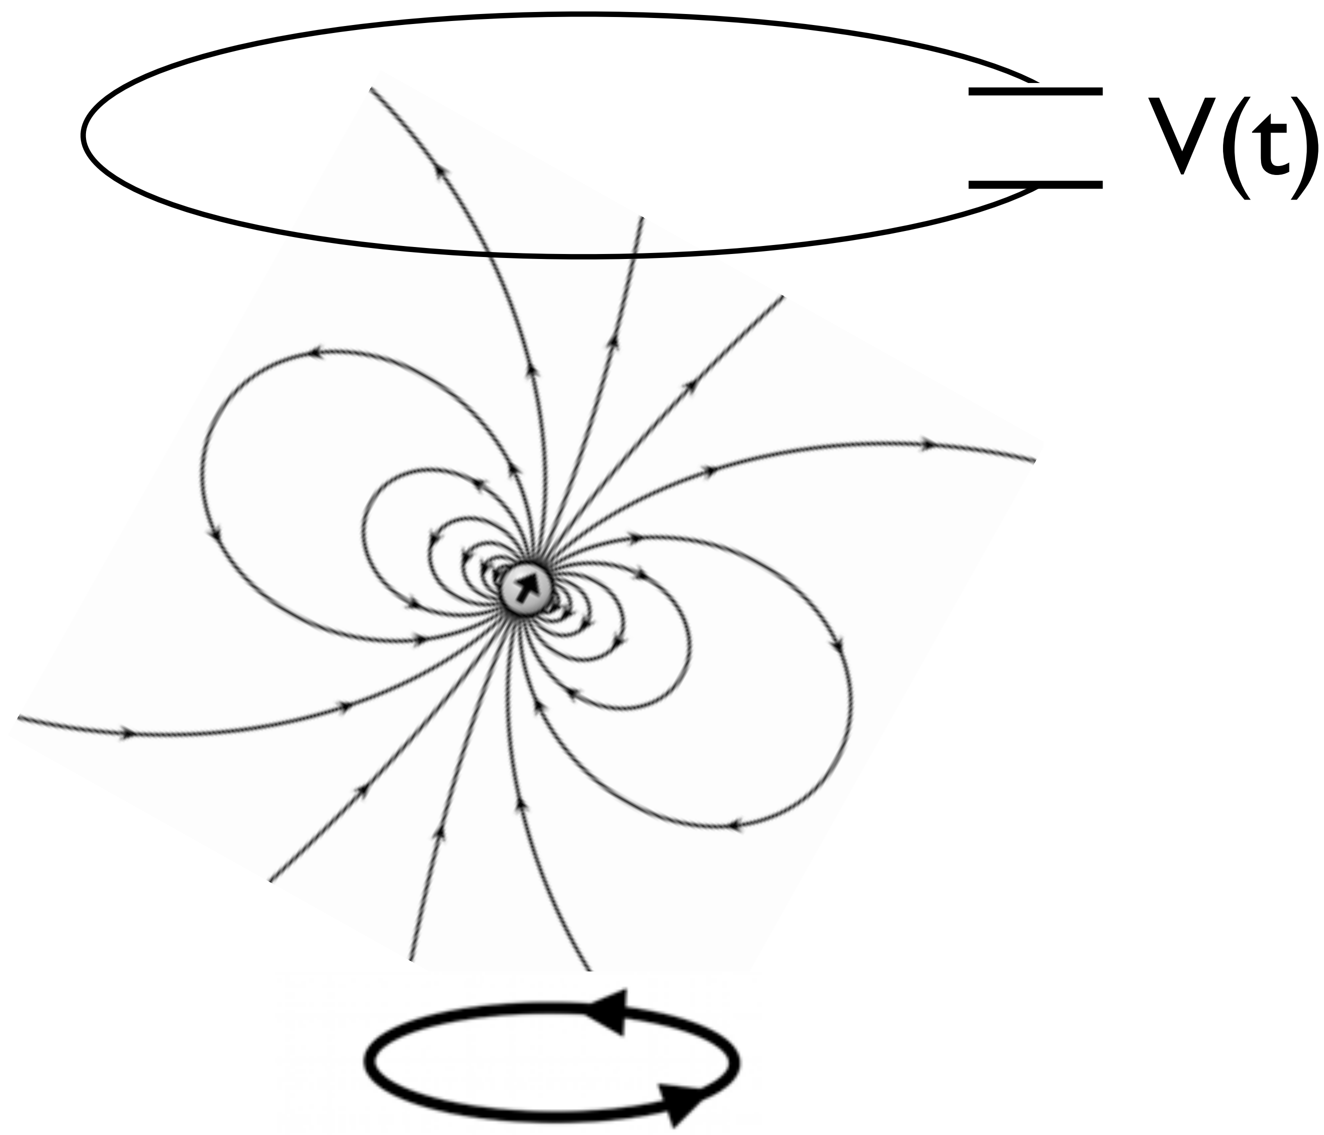

I'd like to skip ahead a little to mention briefly how we'll end up actually detecting a signal from the spinning protons.

A few definitions and preliminary explanations will be required. The magnetic flux is the amount of magnetic field passing through a particular surface (here it will be the surface defined by the loop of wire in the right-hand diagram). What magnetic flux are we referring to here? Well, remember that the protons (like the compass needle) are themselves little magnets. So they induce a magnetic field around them. The flux resulting from that magnetic field passing through the loop is what we're concerned with. In the right hand diagram, the lines of magnetic field **created by the spinning proton** are shown around the proton (the little sphere with an arrow indicating the direction of its magnetic moment).

To fully appreciate why we care about this magnetic flux, we have to mention Faraday's law of induction, which states that the voltage at the terminals of such a loop is equal to the time derivative of the magnetic flux through the surface formed by that loop. In other words, if the magnetic flux through the loop changes, we will observe a time-varying voltage.

In this simplified view, we detect a signal by placing a loop of wire near the sample in question; when there is a time-varying magnetic flux through that loop, we can observe a changing voltage at the terminals of the loop. This time-varying voltage is ultimately what will result in the reconstruction of an MR image. **The precession of the protons is what causes the time-varying magnetic flux through our loop of wire. The time-varying magnetic flux in turn induces a time-varying voltage (the "signal") at the terminals of our loop of wire. We can record this signal and use it to reconstruct an image, as we'll see in the later part of this talk.**

Creating an MR signal: Excitation

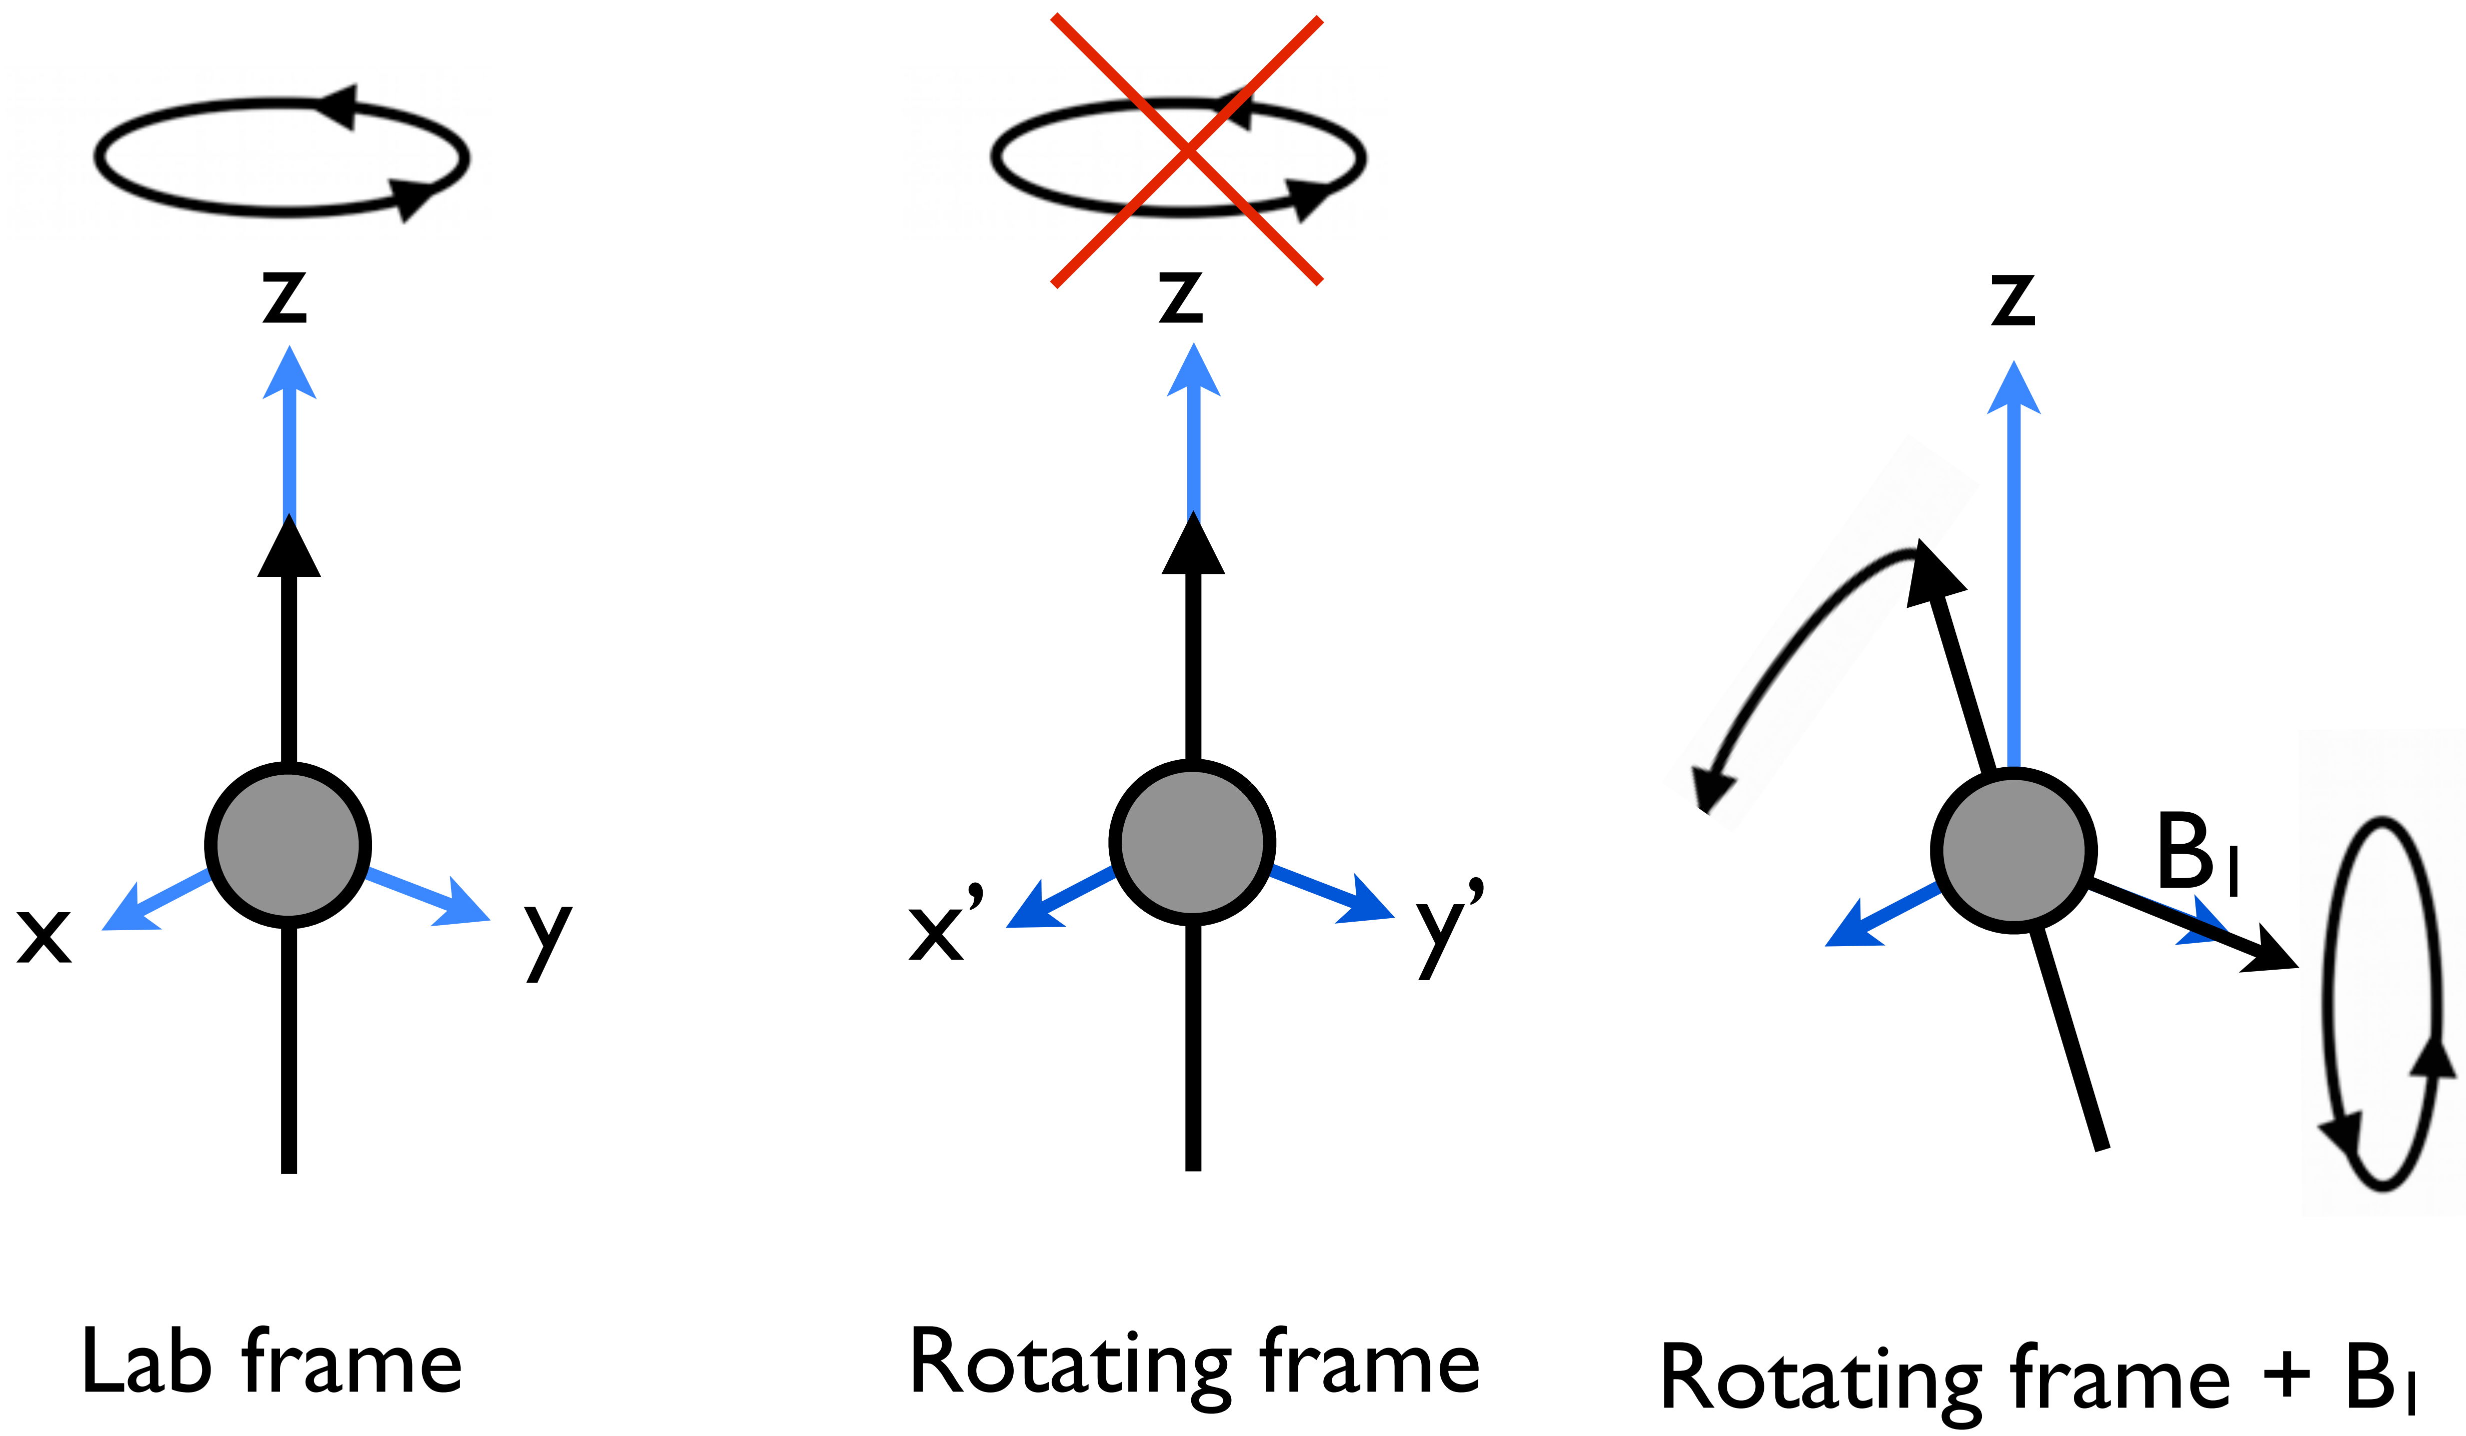

We want to achieve precession around the axis of the magnetic field in order to get a time-varying flux through our signal detection loop. To do that, we need to get the net magnetic moment "off-axis" first, or tip it down from the direction along the B0 field.

Here we'll work in two frames of reference, or two points of view: the lab frame, where we're sitting, and the rotating frame, which is the frame rotating with the protons at their precession frequency. Imagine sitting on a carousel (http://www.google.com/search?q=carousel&tbm=isch): the person on the horse next to you looks stationary to you even when the carousel is spinning around, because they're rotating at the same frequency. This is the essence of the rotating frame.

This means that in the rorating frame, it's as if the main B0 field wasn't there. If it was there, you would detect a precession. But since there's no precession, it's as if B0 was absent (the math for this is much more rigorous, but unfortunately loooong, so we'll skip it). The **effective** magnetic field is zero in the rotating frame.

Now imagine introducing a B1 magnetic field that points along some direction in the x-y plane. But unlike B0, B1 is time varying; furthermore, it's time varying at the same frequency as the precession frequency of the proton; and further, it's spinning around in the same direction (in the sense of clockwise or counterclockwise) as the precessing proton. Well, in the rotating frame, this B1 field will look stationary, and pointed along some direction in the x-y plane (I chose y' here).

What will happen to the proton now? Well, we know how protons react in the presence of a static magnetic field: they act like spinning tops, and precess around the magnetic field. Except this time, the precession will be around B1. This means the magnetic moment of the proton will start to tip down and away from the z axis. If we turn off the B1 field then, the proton goes back to precessing around the z axis, and we'll be able to detect a signal!

Creating an MR signal: Excitation

B_1 field has to be rotating at same frequency as the nuclei, or they will not "see" it

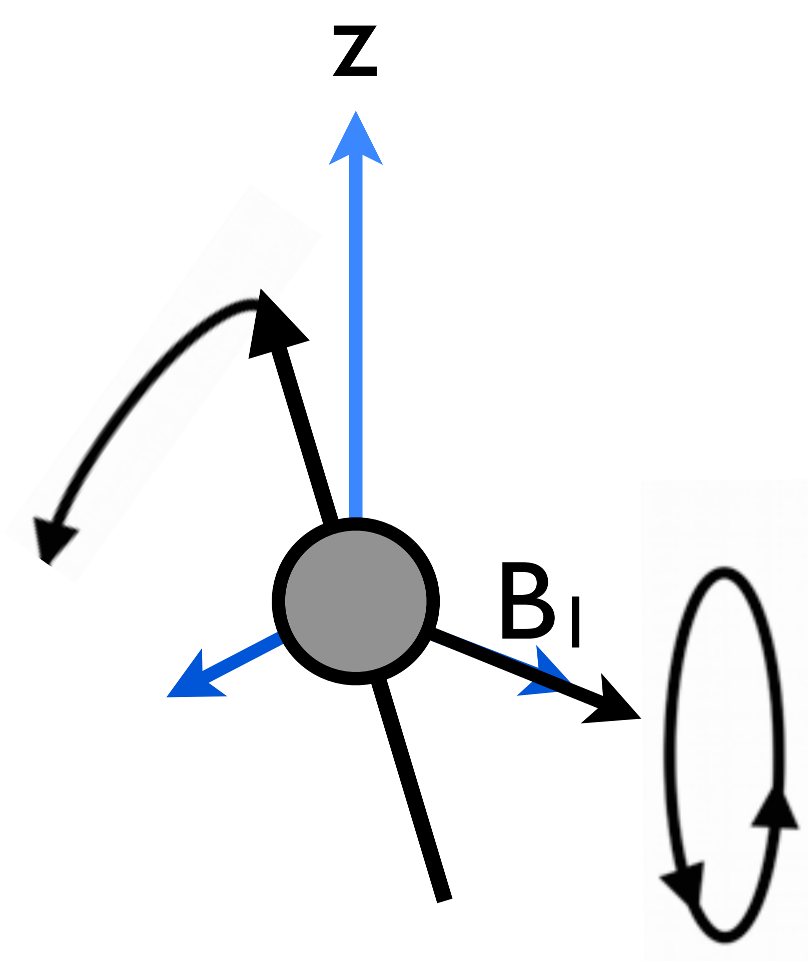

Rotating frame + B_1

After the excitation (90°) 1: precession

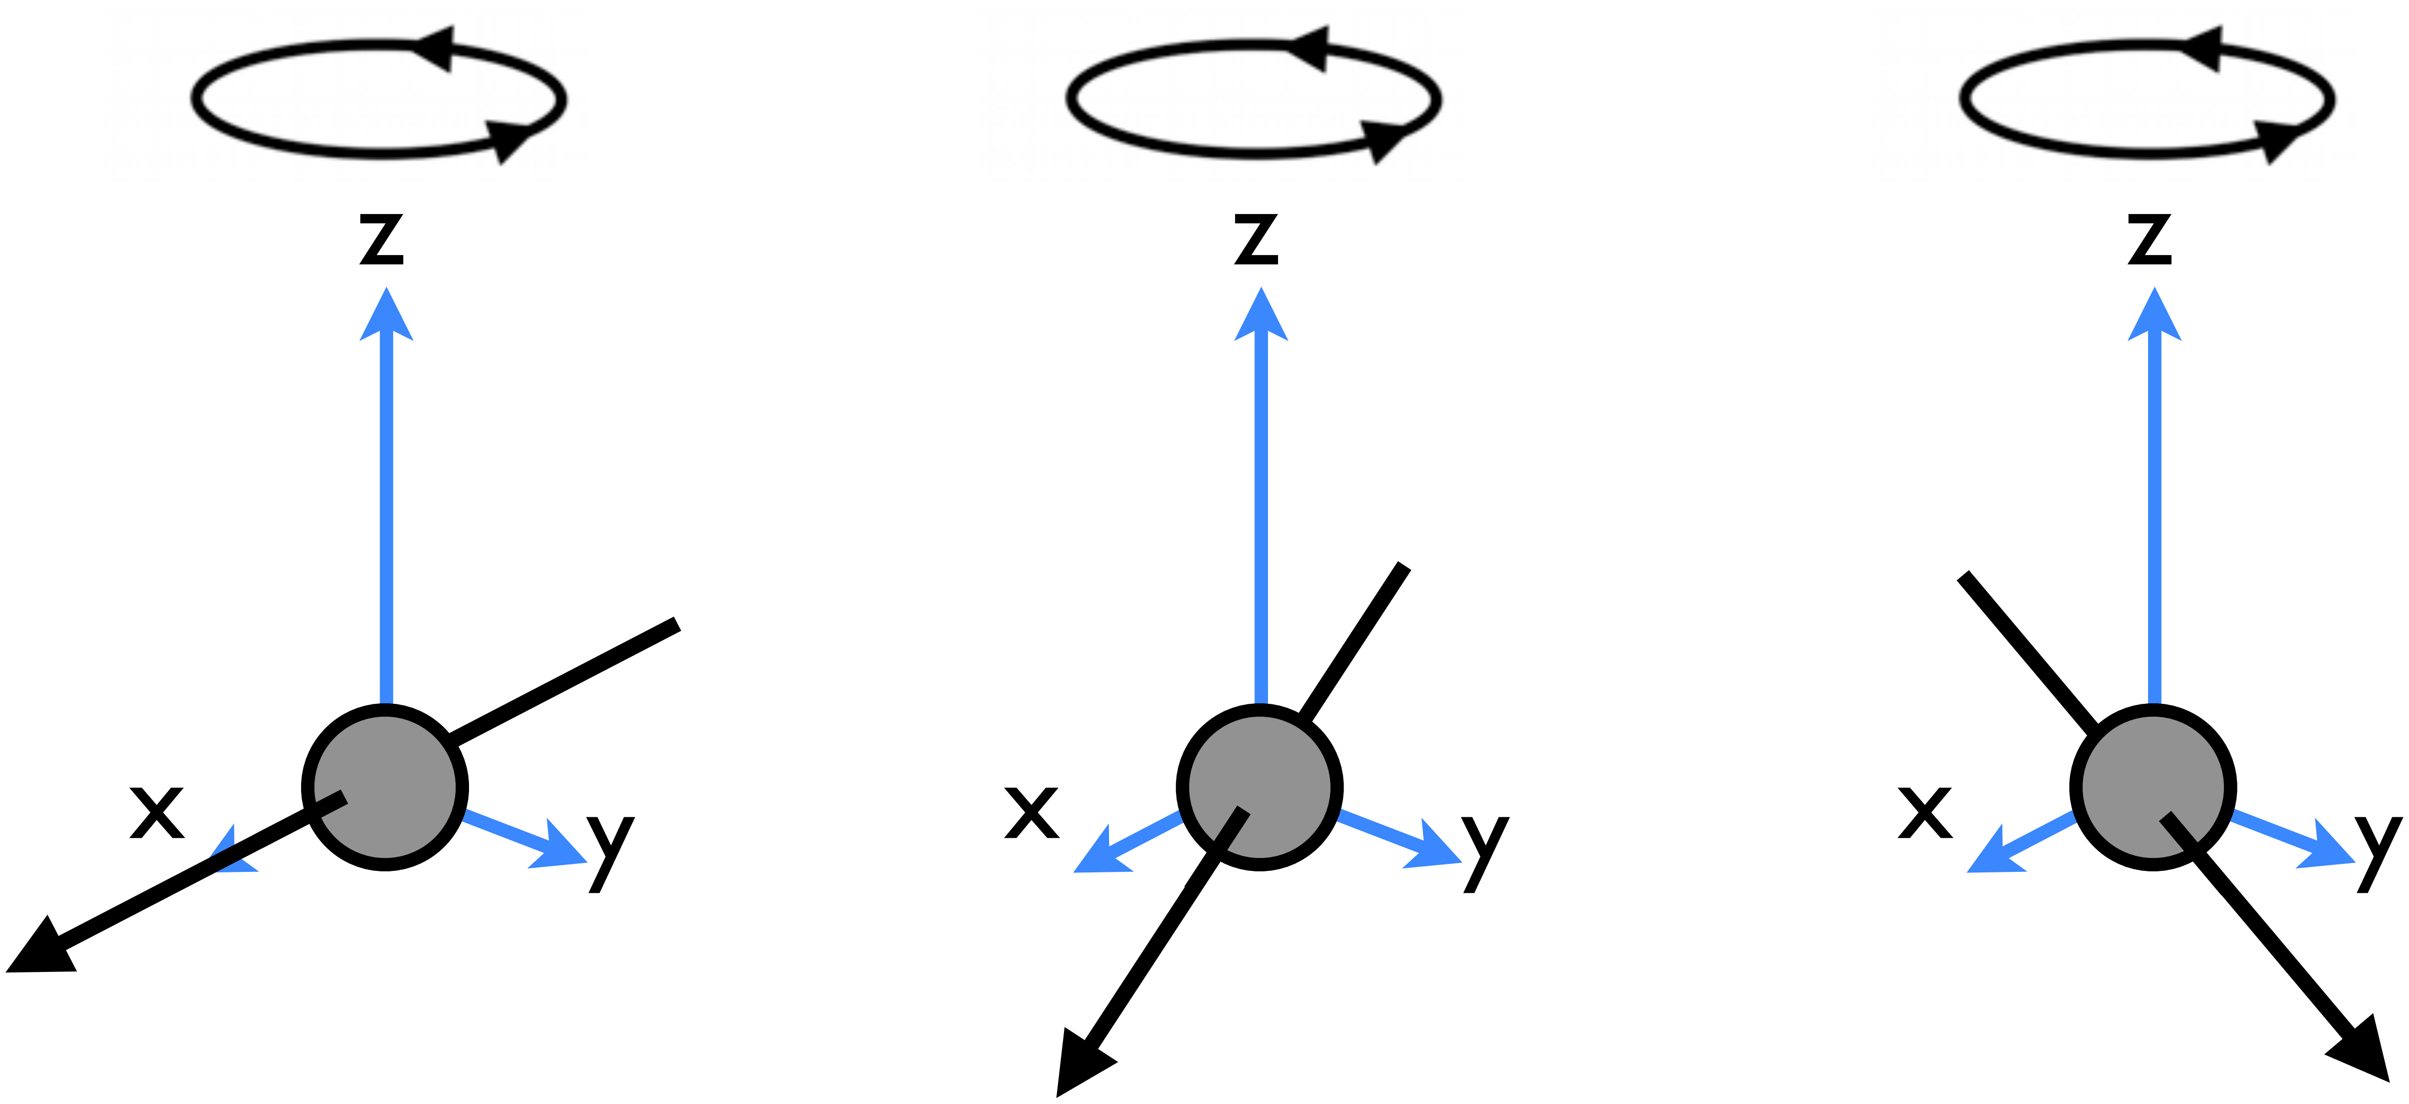

All lab frame

Now we consider what happens after the excitation (tipping away from z). We'll consider 3 phenomena in turn, keeping in mind that they really happen all at the same time, and the separation here is for illustrative purposes.

First, we consider precession: this is familiar from earlier discussion. The magnetic moment will precess around the magnetic field axis, at 42.58 MHz for a proton in 1 T field.

After the excitation 2: dephasing (T2*)

T2* due to *static* differences in effective B.

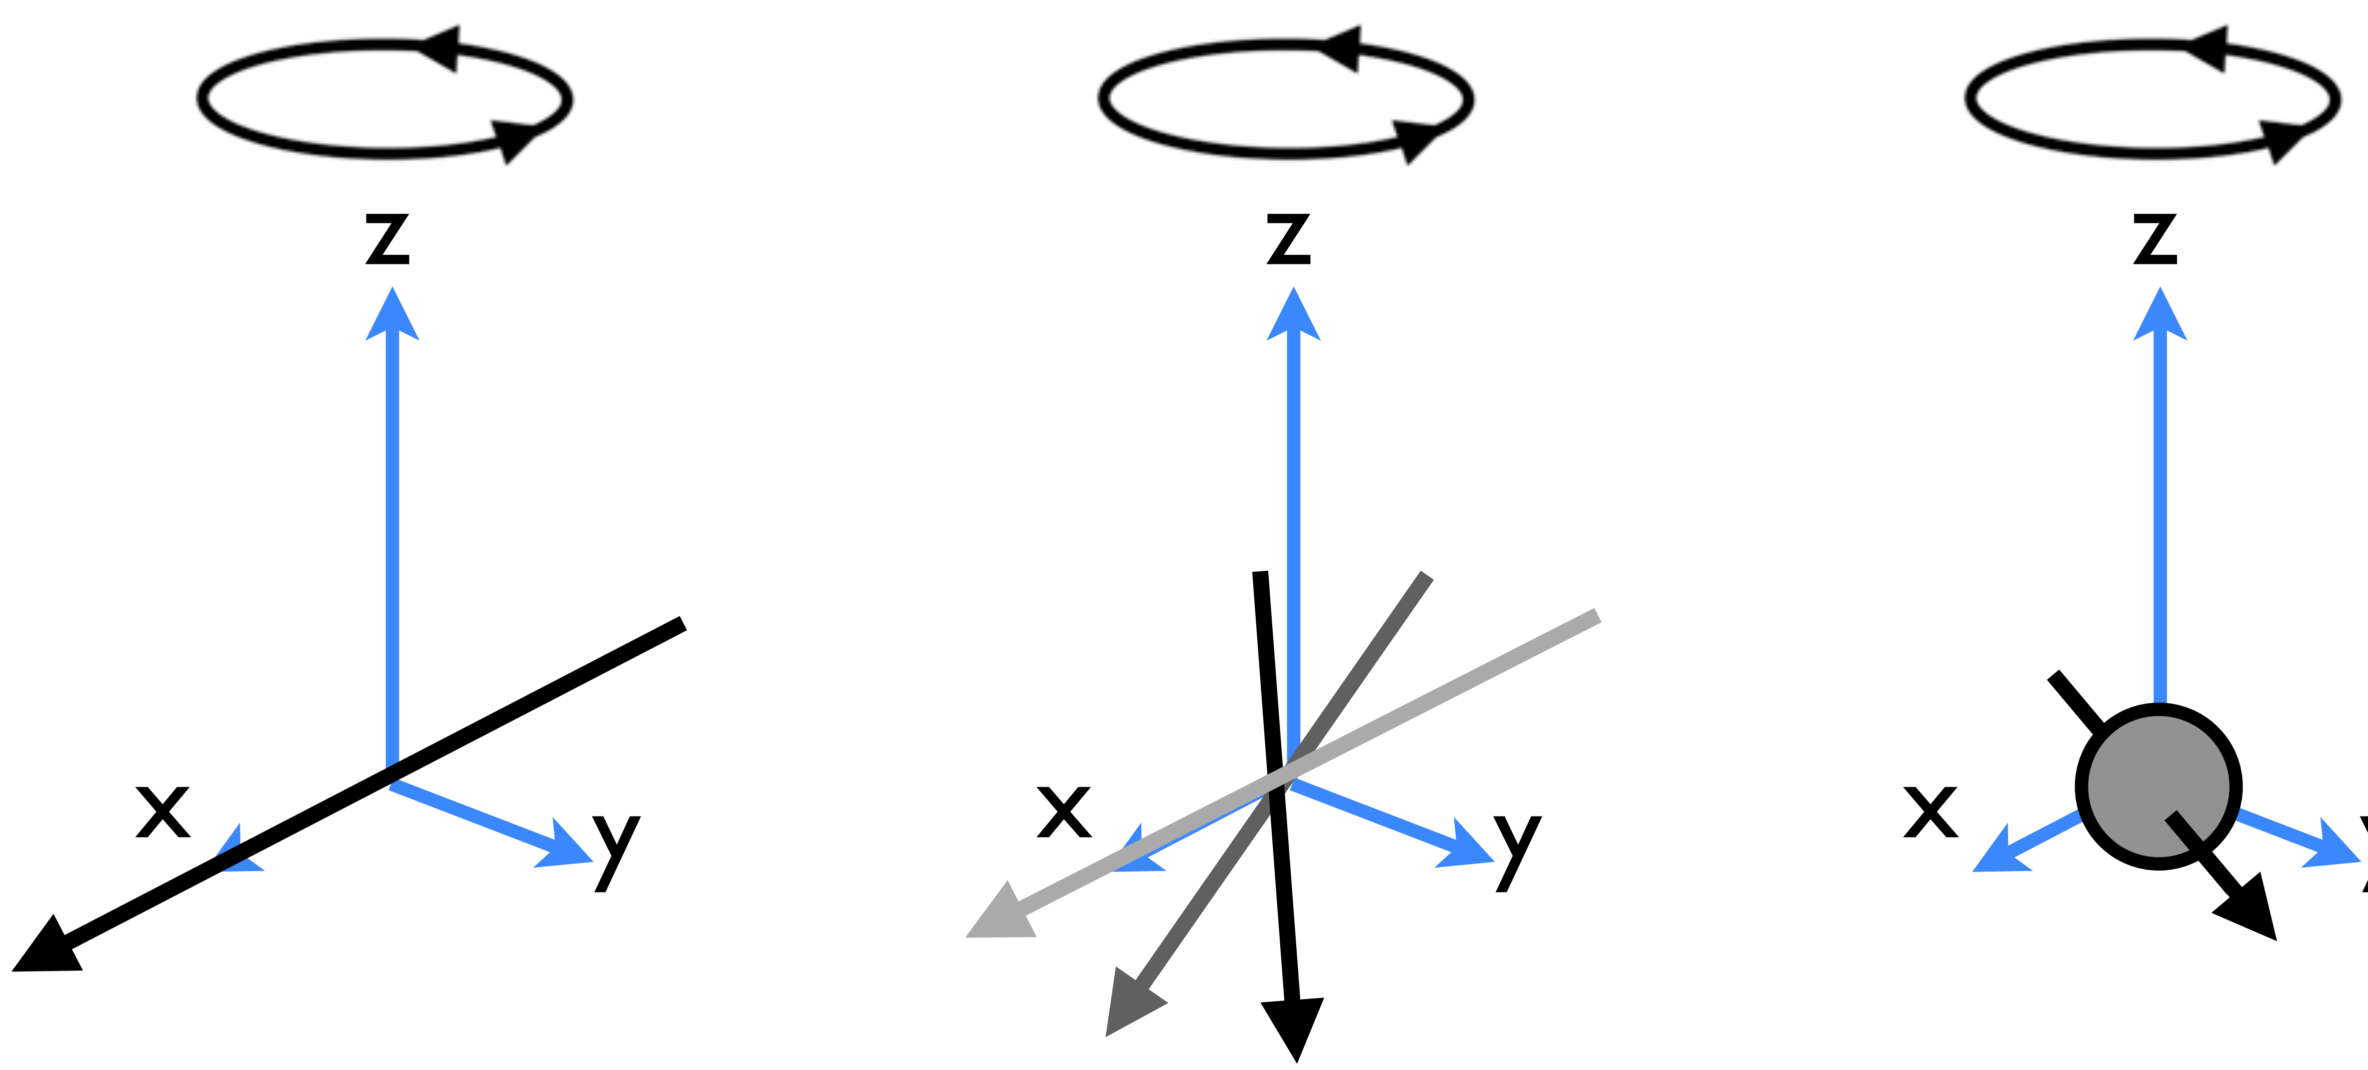

Some nuclei will precess faster than others

Net magnetization is smaller

The second phenomenon we consider is dephasing. Now we think of several nuclei (here, I draw 3 different magnetic moments). Let's imagine that these nuclei are at different positions in our sample. I mentioned early on that the the B0 field is not *perfectly* uniform. That means that nuclei in different places in our sample will experience different strength B0 fields. Furthermore, in a biological sample, these nuclei will be in very different environments (water inside a cell, outside a cell, near a sinus or near iron deposits, etc), all of which further contribute to the fact that they experience a different total magnetic field.

The main point is this: different magnetic field strength ---> different frequency of precession. So some nuclei will "get ahead" of others as they precess, while others will fall behind.

In the first digram, all three nuclei *just* started precessing after the excitation. They are all at the same position. In the second diagram, the nuclei that experience a higher total B field have moved ahead of the pack. In the third diagram, I show the vectorial sum of all the precessing nuclei. Over time, as they spread apart, the net magnetic moment gets smaller, as the "slow" and "fast" nuclei partially cancel each other out (they point along different directions).

Eventually, all the magnetization in the x-y plane is gone, and we can no longer get a signal at our detection loop.

After excitation 3: return to z (T1)

Rotating frame

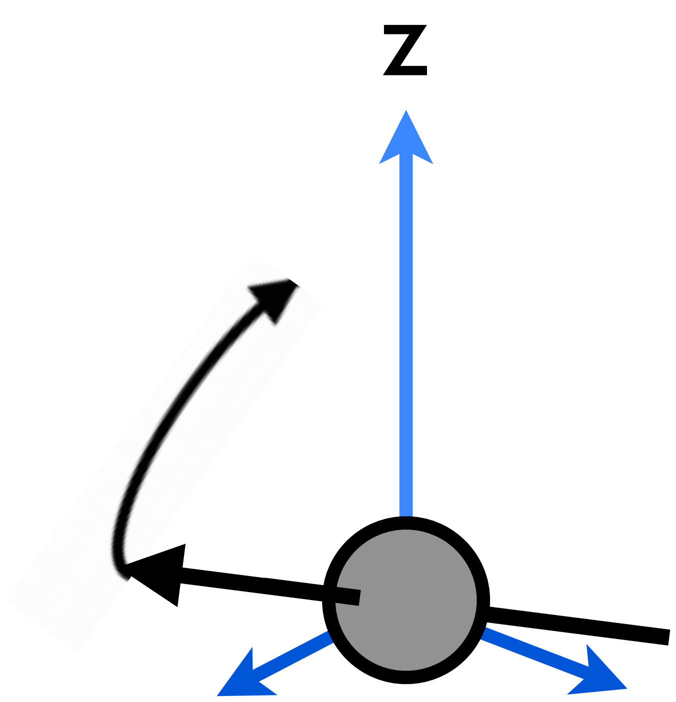

Technically, Mxy goes to 0, and Mz goes back to original value

e.g. what about 180° flip?

The third phenomenon we consider is the return to equilibrium, which is a small net magnetization pointed along z.

The picture here is a little more difficult to draw, and this diagram is technically wrong, but I use it to remind you that the magnetization returns to pointing along the z direction. In actuality, two things are happening at the same time: the magnetization in the x-y plane (Mxy) is going to zero due to dephasing, and the magnetization in the z direction (Mz) is little by little returning to its initial value. The latter happens as nuclei exhange energy with the environment and one another, which changes their magnetization direction, with a slightly preference for being along z.

After excitation 3: return to z (T1)

Rotating frame

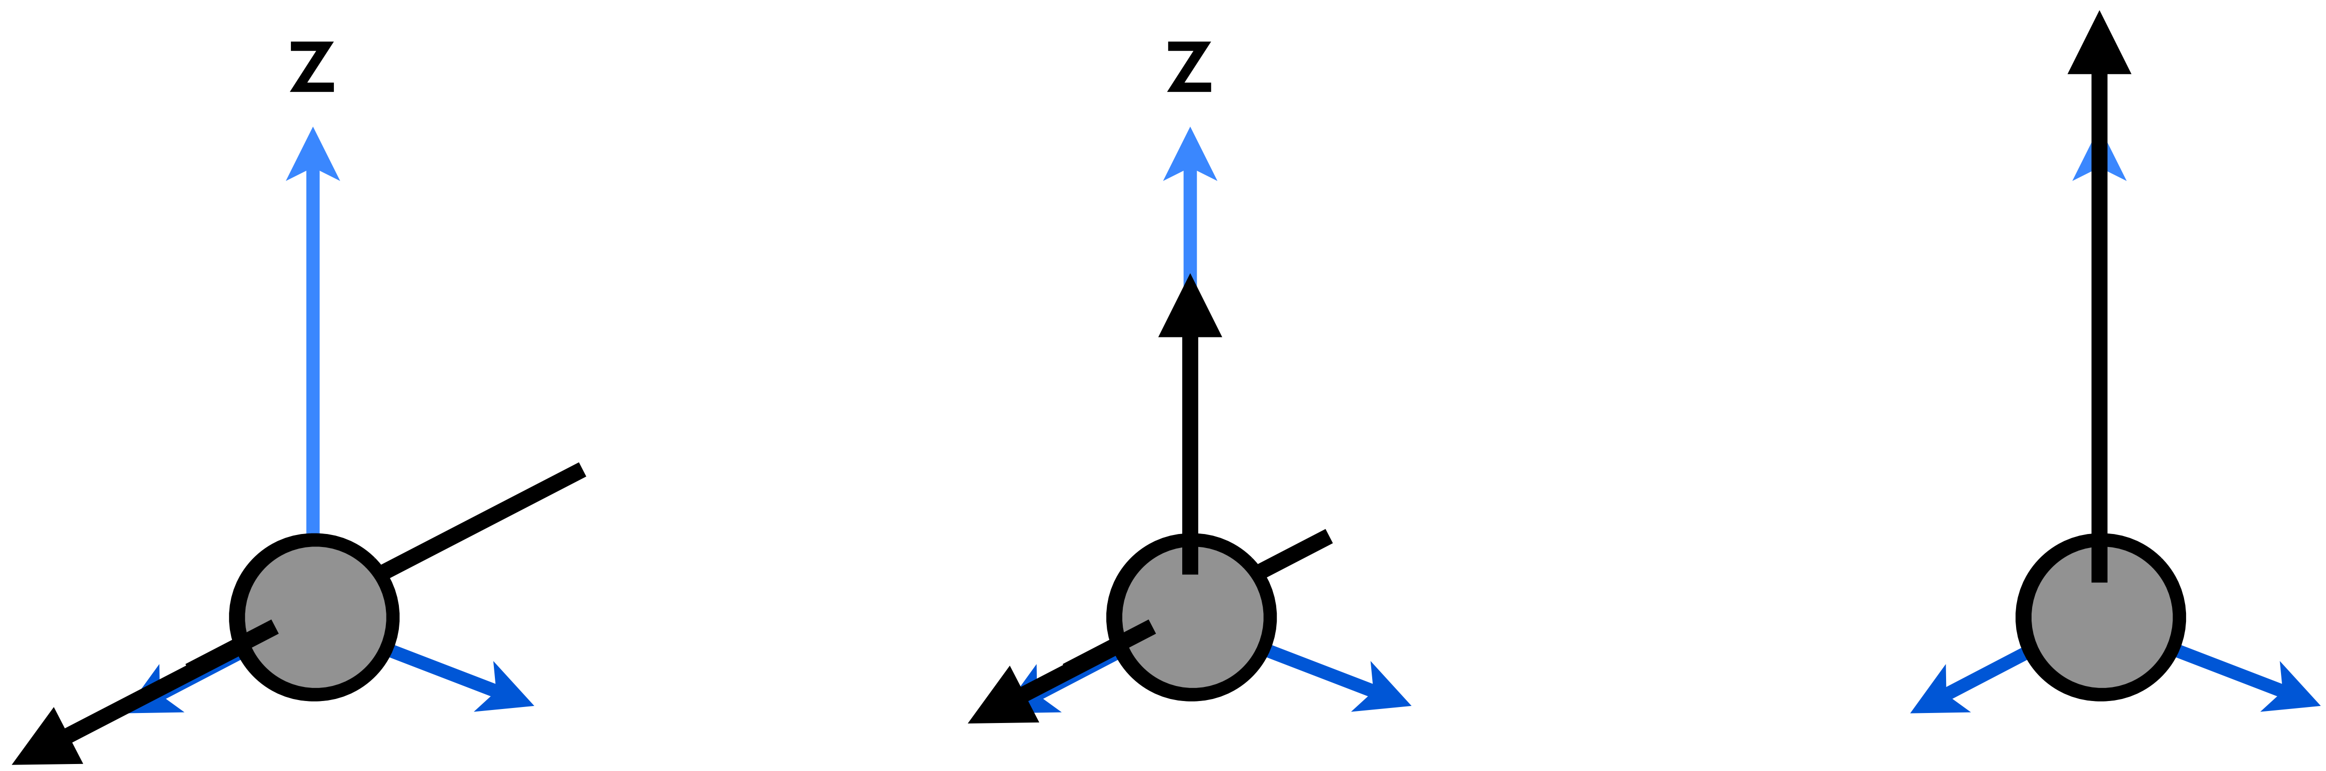

This diagram, in the rotating frame, shows more accurately the processes of dephasing and return to equilibrium.

The two vectors shown are Mz and Mxy. In the first diagram, Mxy is at its maximum, and Mz is zero. In the second diagram, Mxy has shrunk due to dephasing, while Mz is starting to return to its equilibrium value. In the third diagram, Mxy is zero, and Mz is back to its initial value.

Detecting the signal & contrast

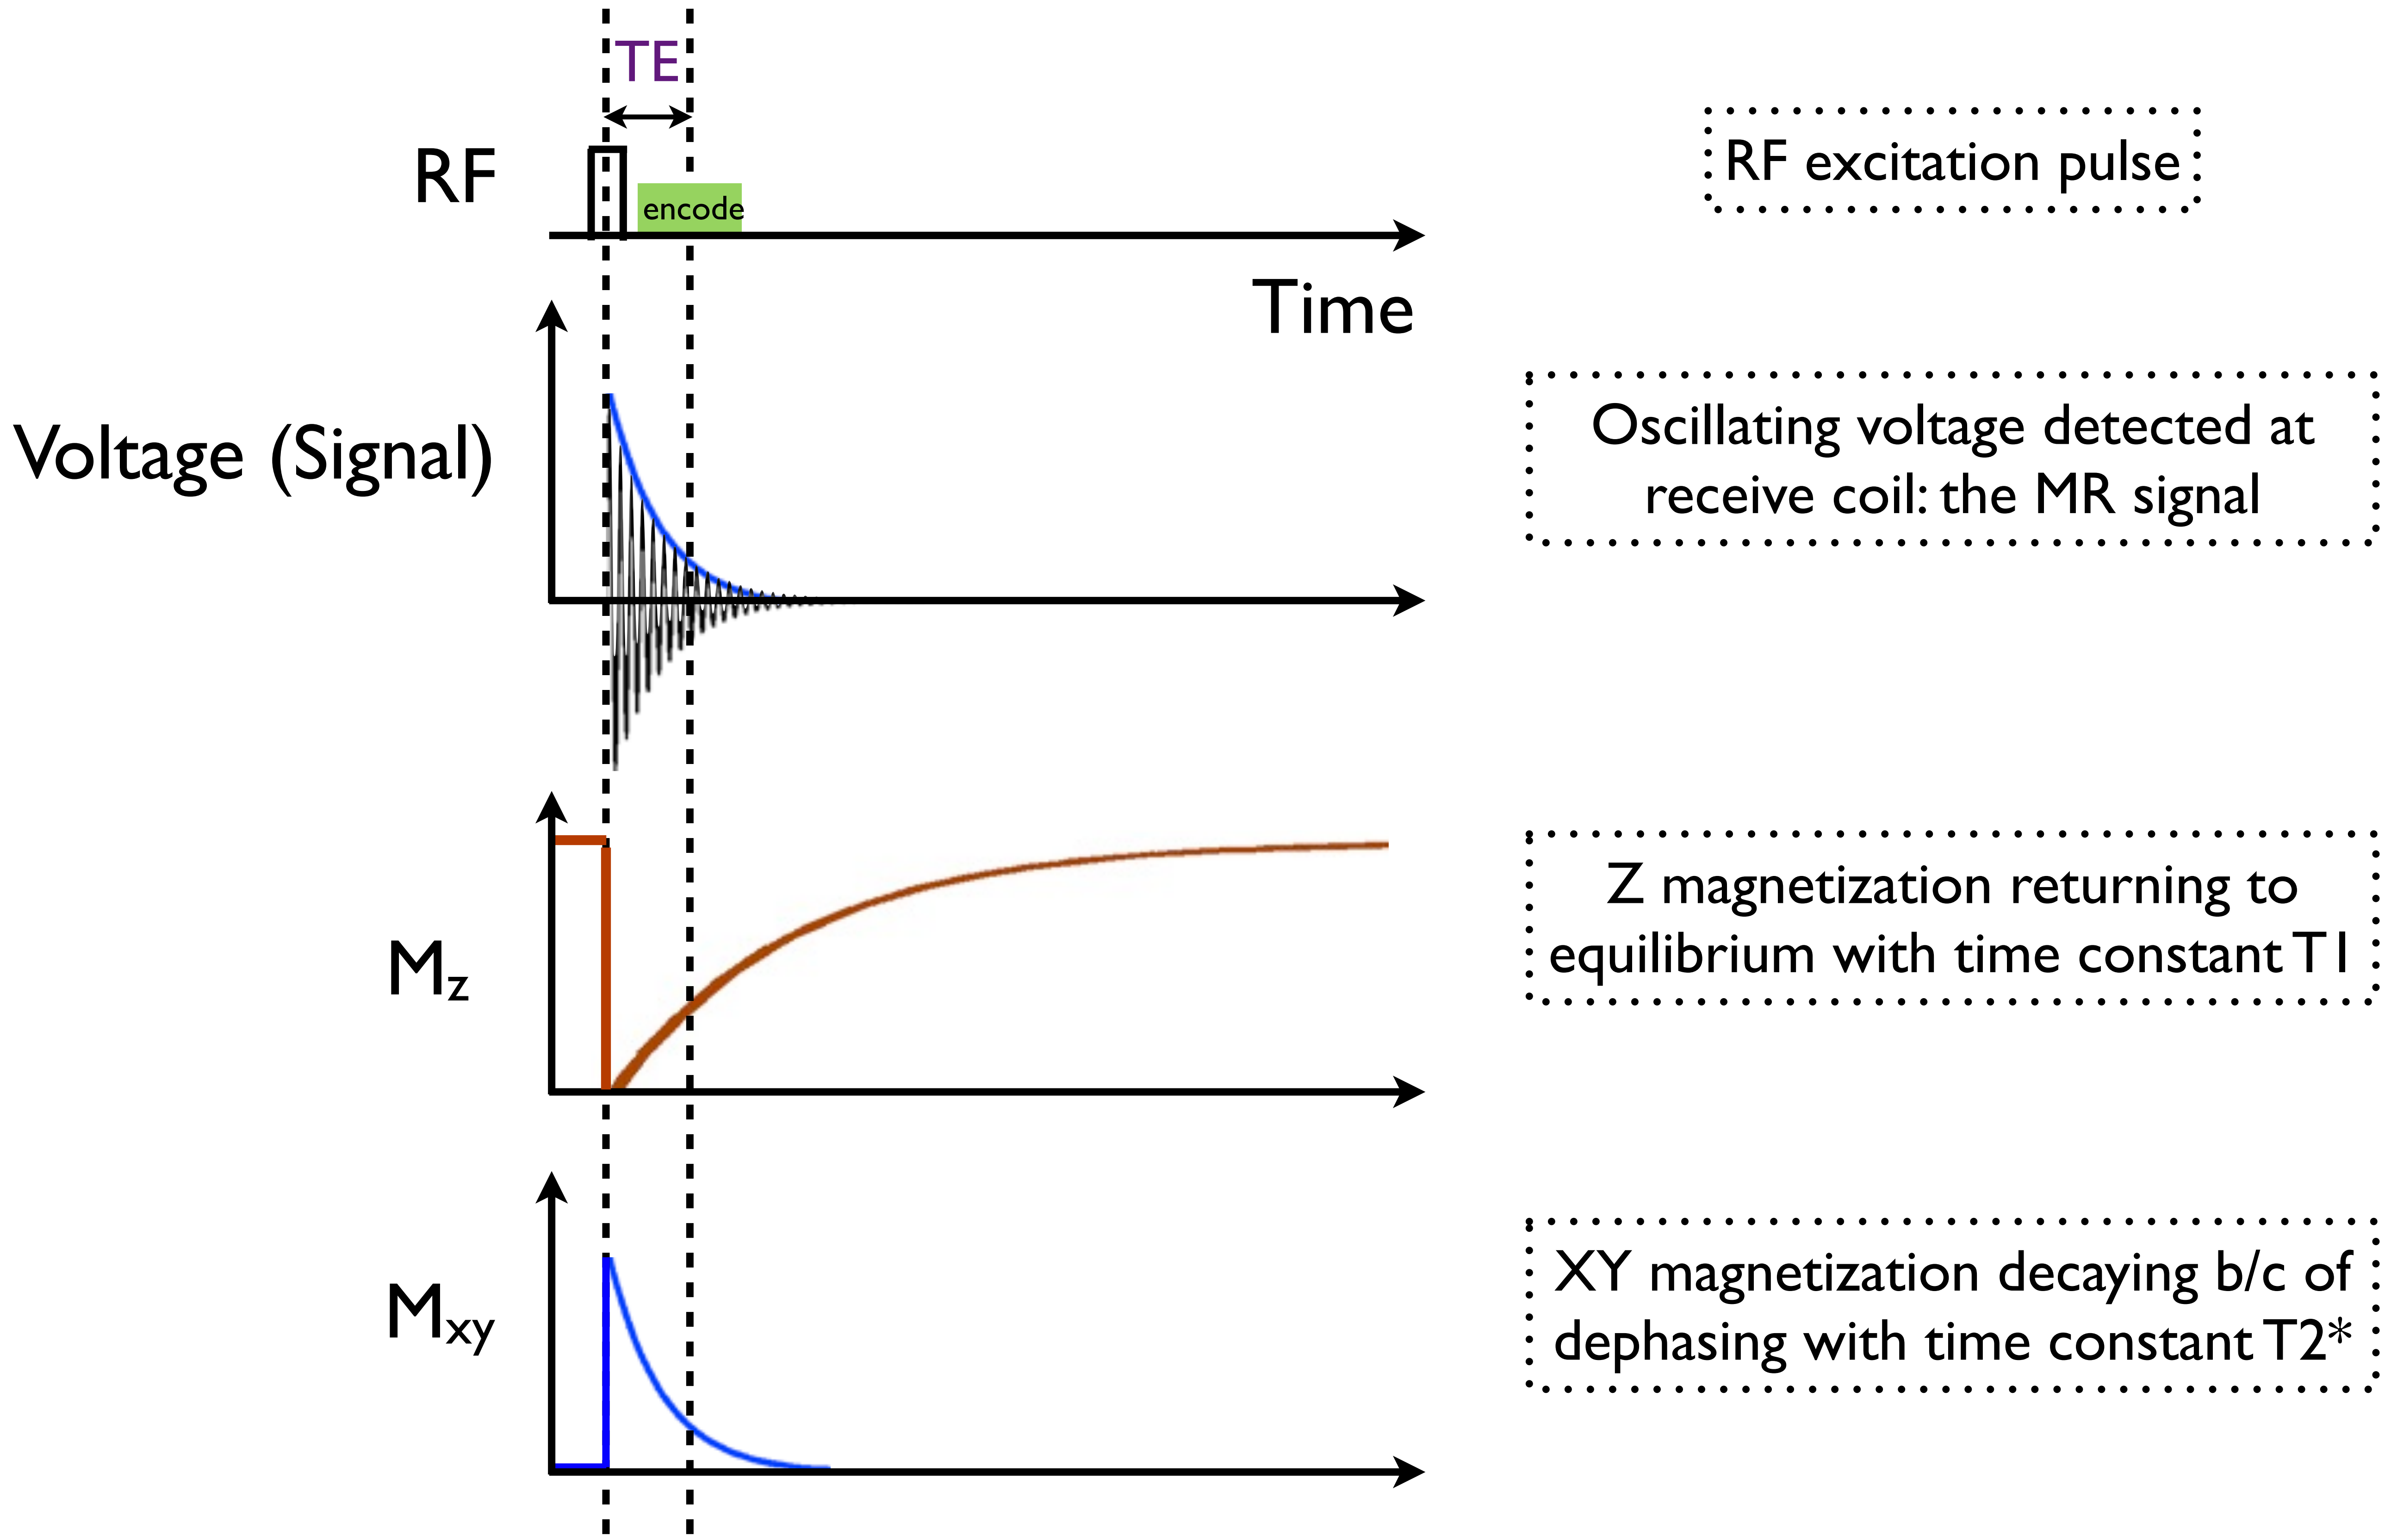

Now that we have all the pieces together to understand how a signal is created, we can look at the time course of the quantities we've described to understand the experiment as a whole.

This slide shows the result of applying *one* RF excitation. In an MR experiment, we repeat this sequence many times over, but we'll begin by analyzing one such excitation.

At the top, I am showing a rectangular RF excitation pulse. This represents the fact that we turn on the B1 field for a length of time corresponding to the width of the rectangle, at a voltage corresponding to the height of that rectangle. The width and height are chosen to as to tip the magnetization down from the Z plane down to the XY plane. We also see a rectangle marked "encode". It is during this time that we actually record a signal. TE is the "time to encode", commonly called "Echo Time", although that's a slightly confusing choice; it is defined as the time between the middle of the excitation pulse to the middle of the encoding. We can choose to place the encoding closer or further away from the excitation, and that will result in different contrasts being emphasized; we'll discuss that in the coming slides.

We see in the second plot that an MR signal is detected as soon as the magnetization tips away from the Z axis, and decays as dephasing occurs.

In the third plot, we see that the Z magnetization returns to its equilibrium value exponentially, with a time constant T1, but that take longer than the dephasing. In other words, even after the signal is gone, we still have not returned to the equilibrium condition.

Finally, in the last plot, we see the XY magnetization. This is zero at first (net magnetization in the Z direction at equilibrium), then quickly reaches a maximum as we tip the magnetization away from Z by 90 degrees, and then it decays due to dephasing with time constant T2*.

A couple things to note before we move on:
--T1 and T2(*) are tissue dependent, e.g. grey matter and white matter will have different T1 and T2 rates (for a given field strength, etc)
--TE is a sequence parameter that we can choose
--TR (not shown) is the time between one RF excitation and the next; we also control this parameter
--TE and TR are the parameters we control to modify the contrast

Detecting the signal & contrast

When in this process do we record an image?

- Equilibrium
- RF Excitation
- Precession & dephasing
- Return to equilibrium

Proton density weighting

T2 weighting

T1 weighting

The choice of where we place the encoding will determine the different contrast we get in the image. Here we see the 4 steps I describe in the time course of the experiment, and the different points at which we can encode, resulting in a proton density, T2 or T1 weighting of the image contrast.

Proton density weighting

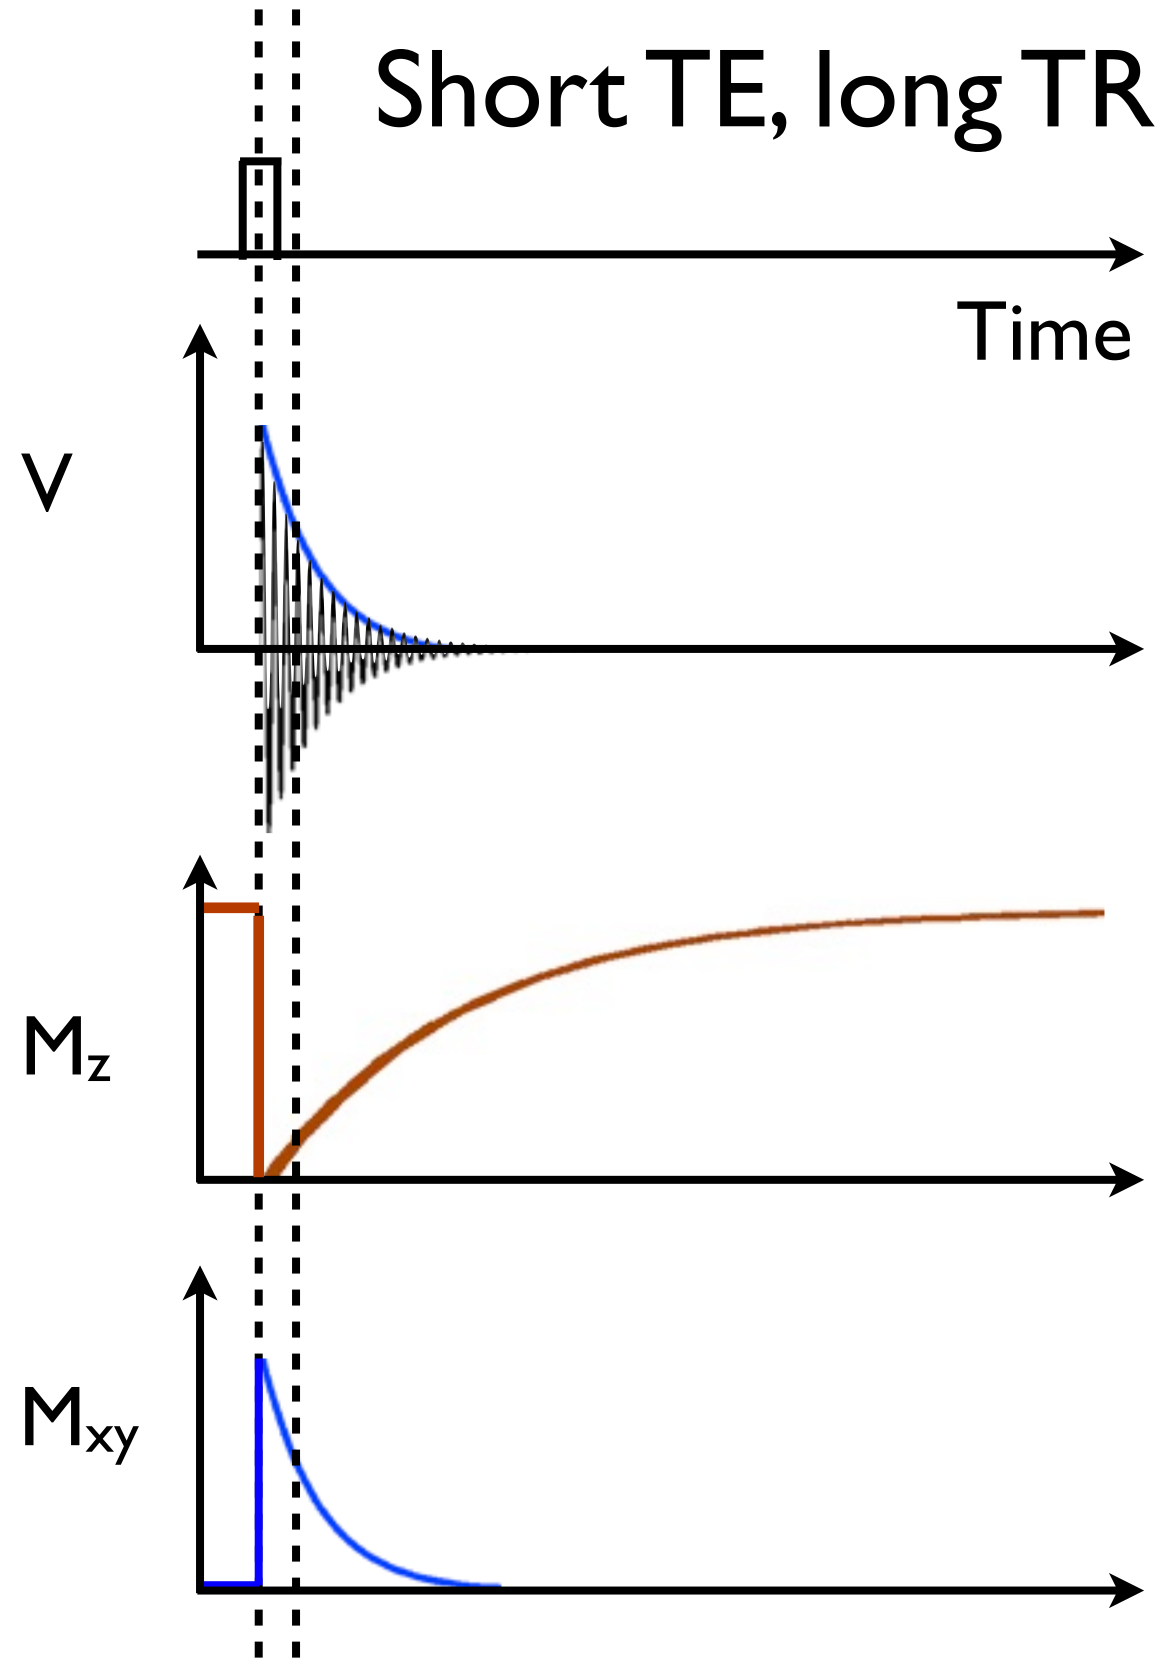

CSF > gray > white

With a short TE and a long TR (not shown), we do not allow tissues with different T2 rates to “differentiate” themselves due to different rates of dephasing. The long TR ensures that all tissues return to initial magnetization before the next excitation, which also removes the possibility of contrast arising from different T1 rates.

The only contrast left is the number of protons in each imaging voxel. The more protons contribute to the signal in a given voxel, the brighter that voxel will be. This is proton density weighting. Note that all MR images are at least slightly weighted by proton density: a voxel with very few or no protons to contribute to the signal will always be dark (e.g. air), regardless of your choice of TE and TR.

T2-weighted spin echo image

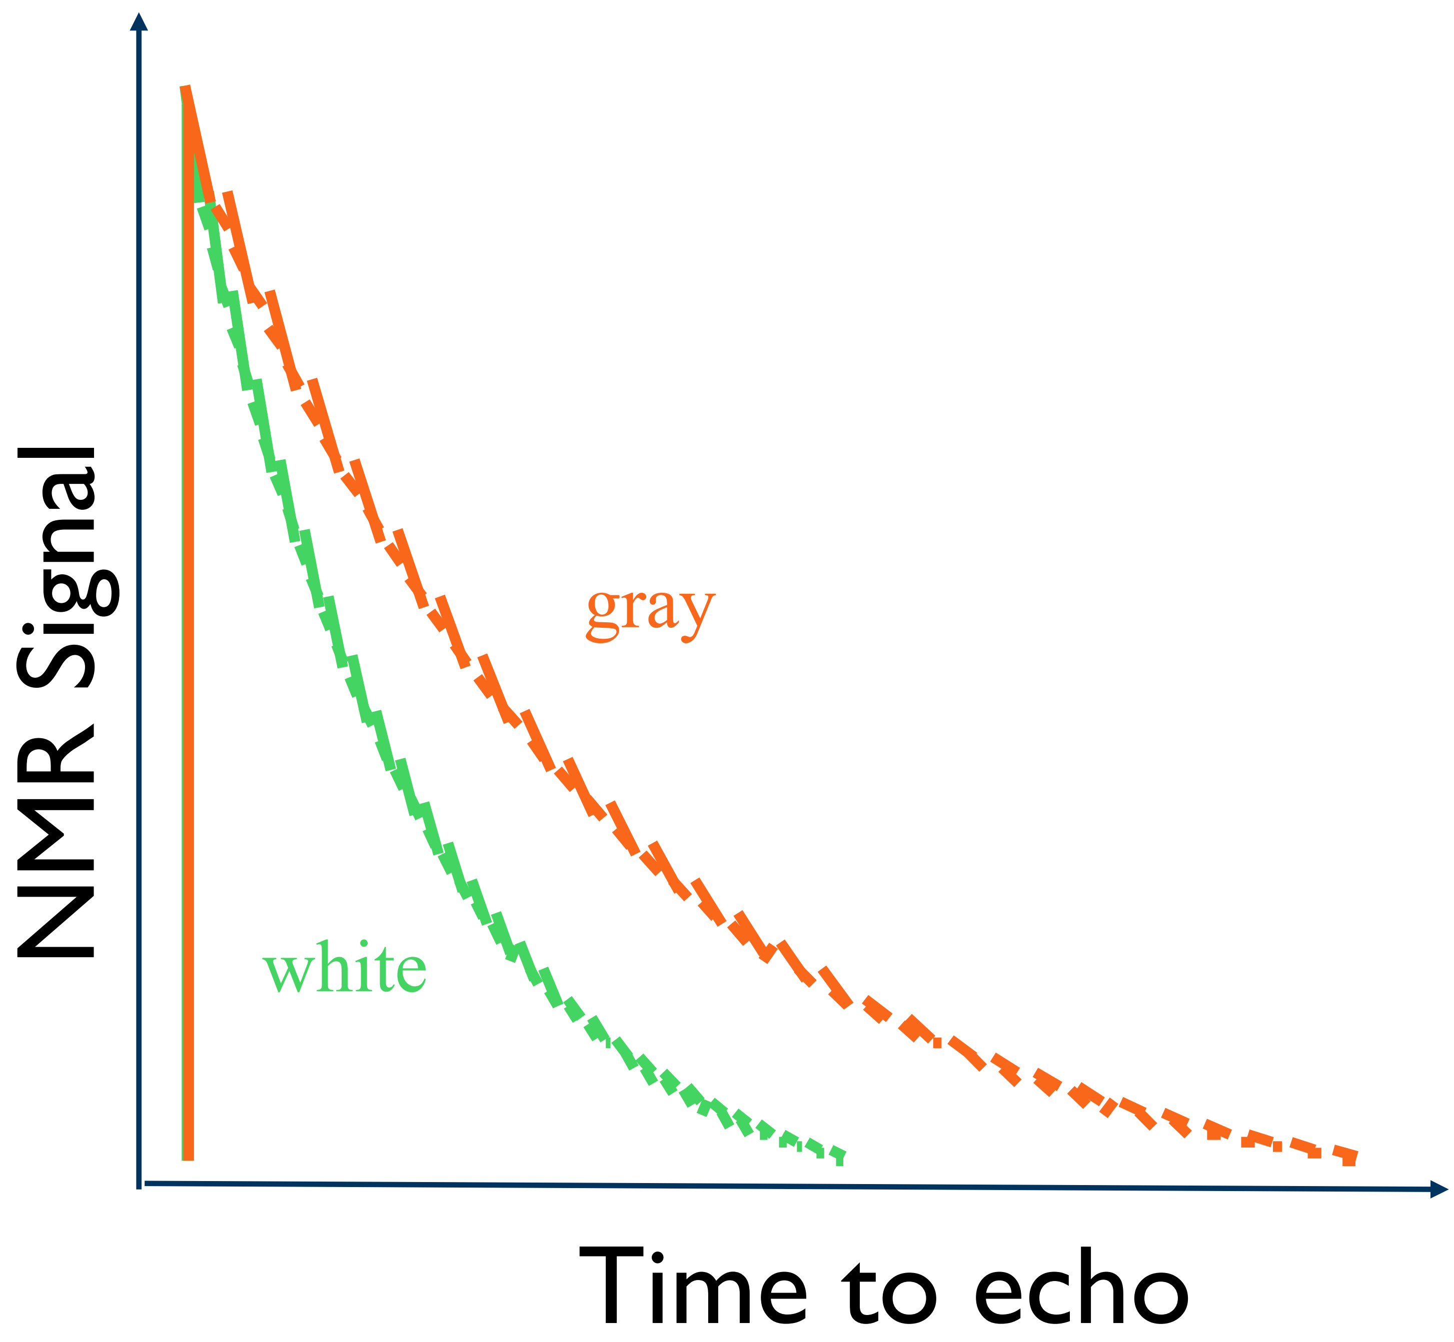

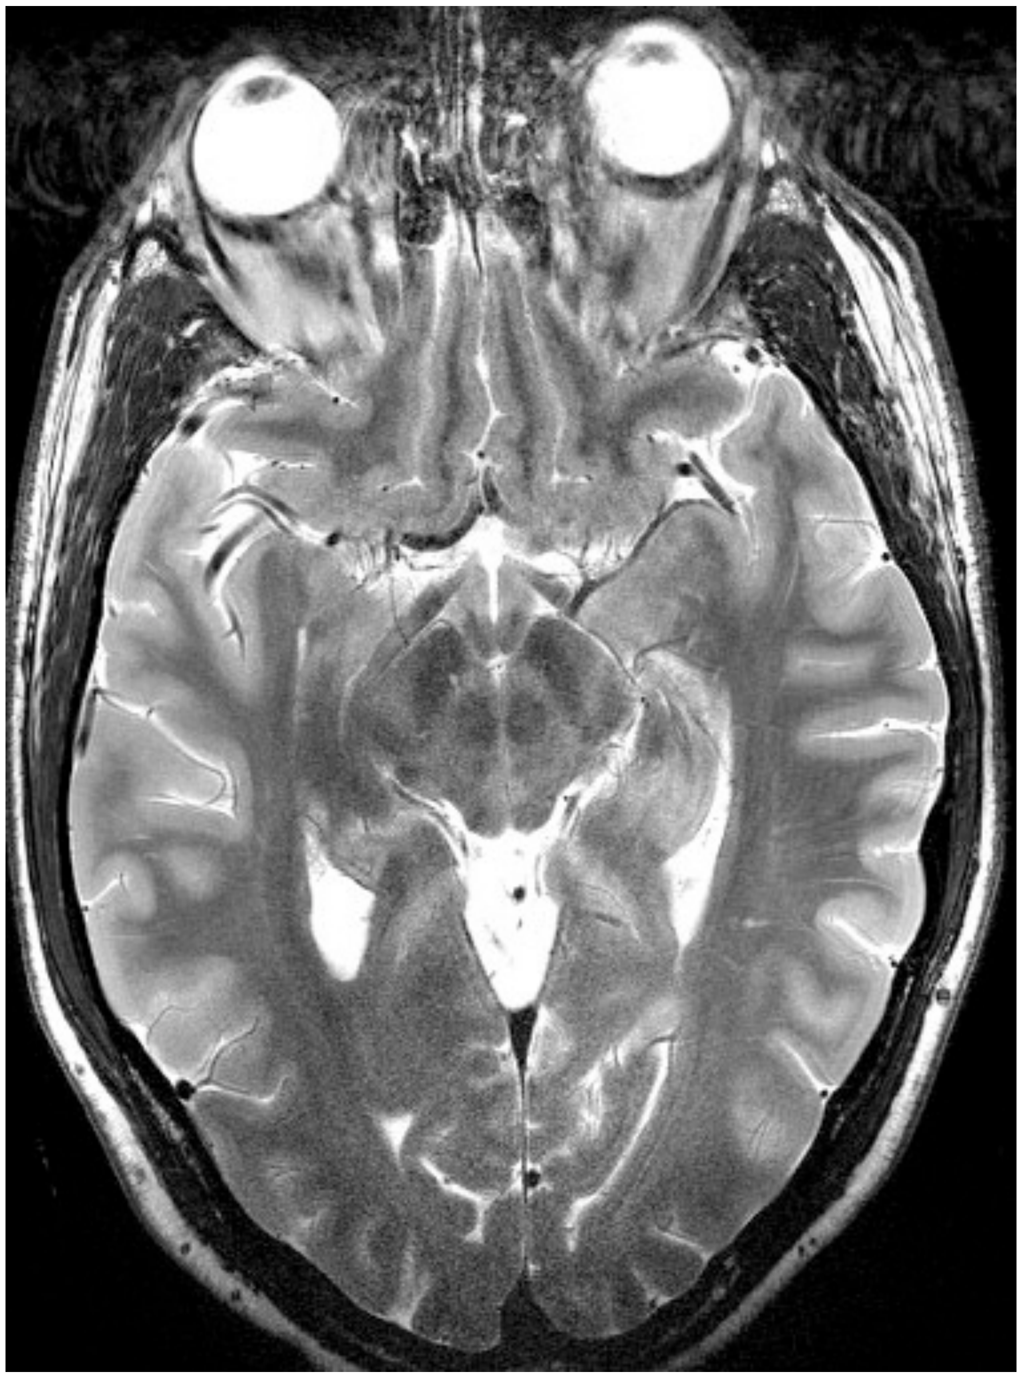

If we choose a relatively long TE (but not so long that the signal has decayed completely), we emphasize the differences in T2* among the different tissues.

T1 weighting

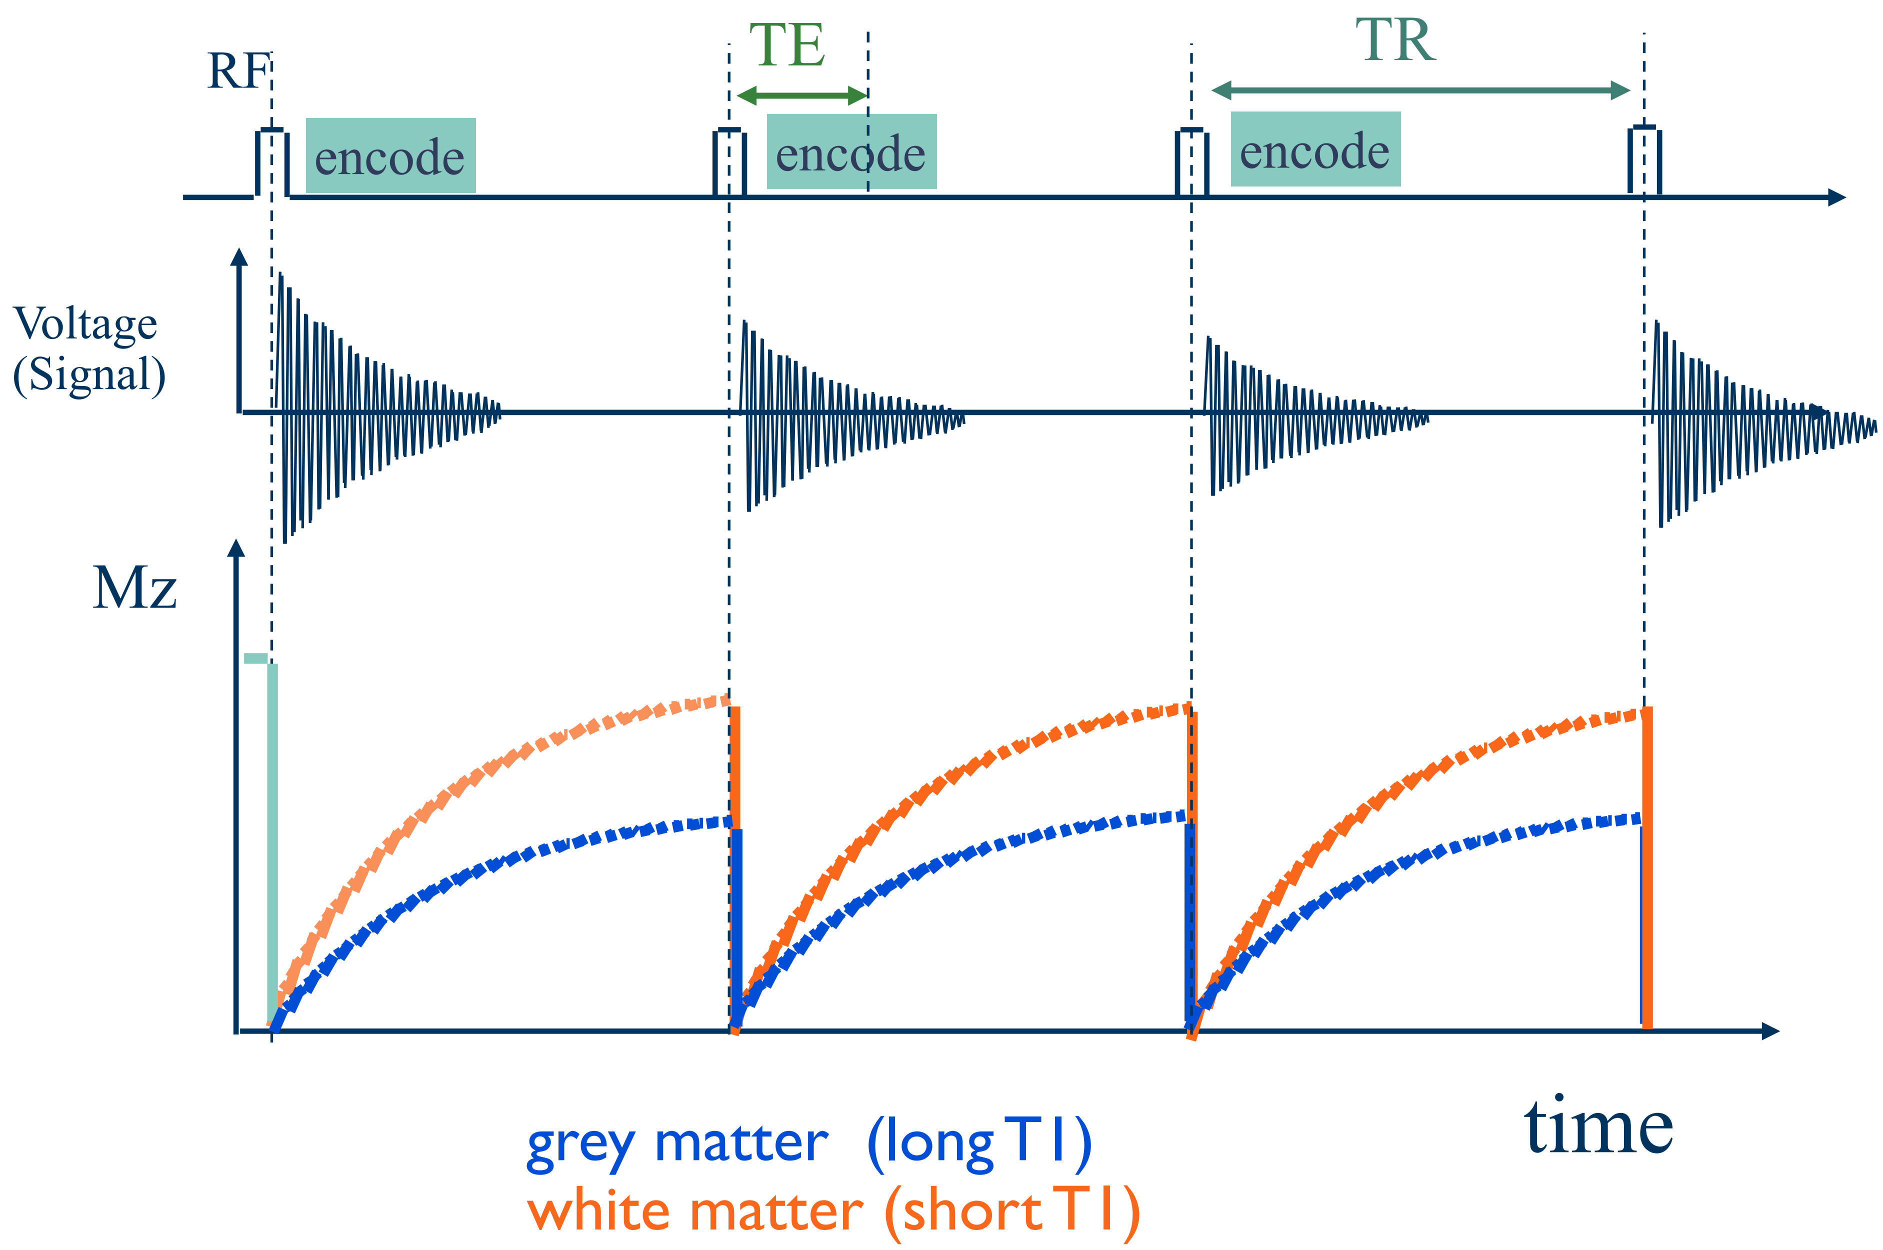

A short TR (and short TE) allows us to emphasize the T1 weighting: tissues whose Z magnetization takes longer to return to equilibrium have few protons “ready to be excited” at the next RF excitation. As a result, fewer protons in that particular voxel will contribute to image brightness.

T1 weighting

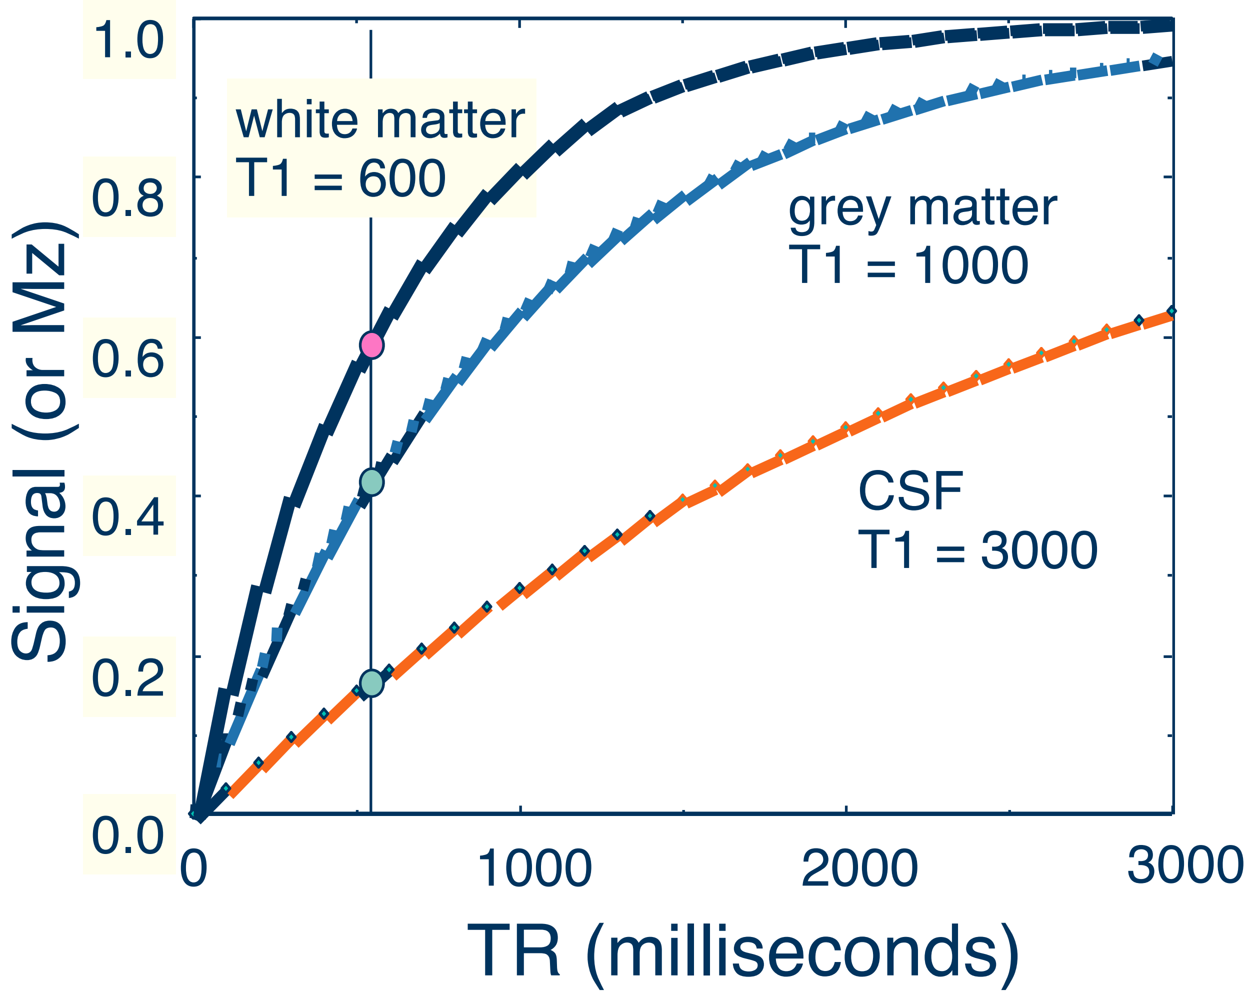

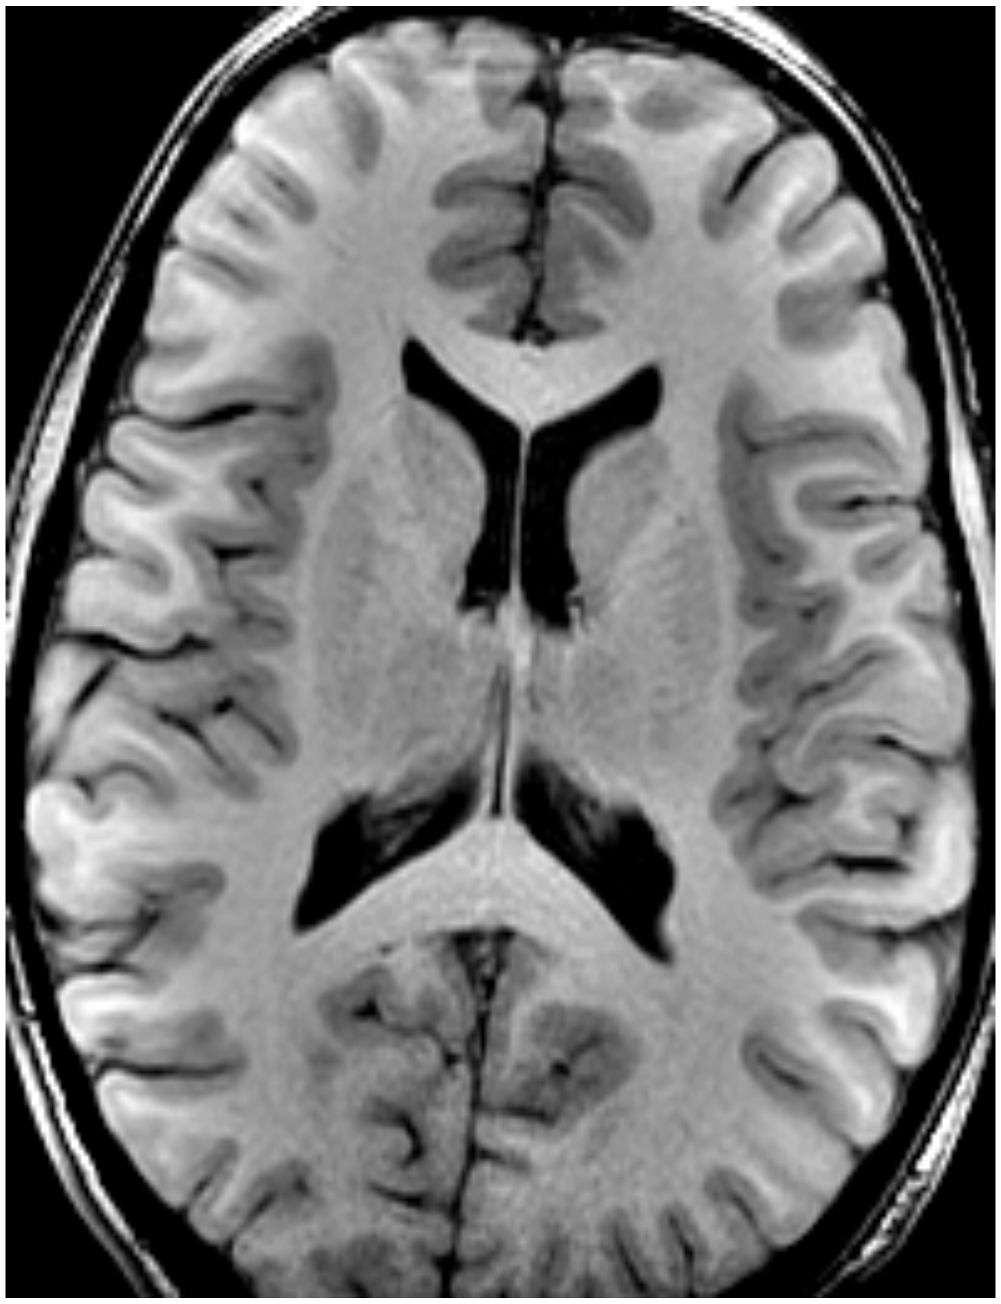

Contrast summary

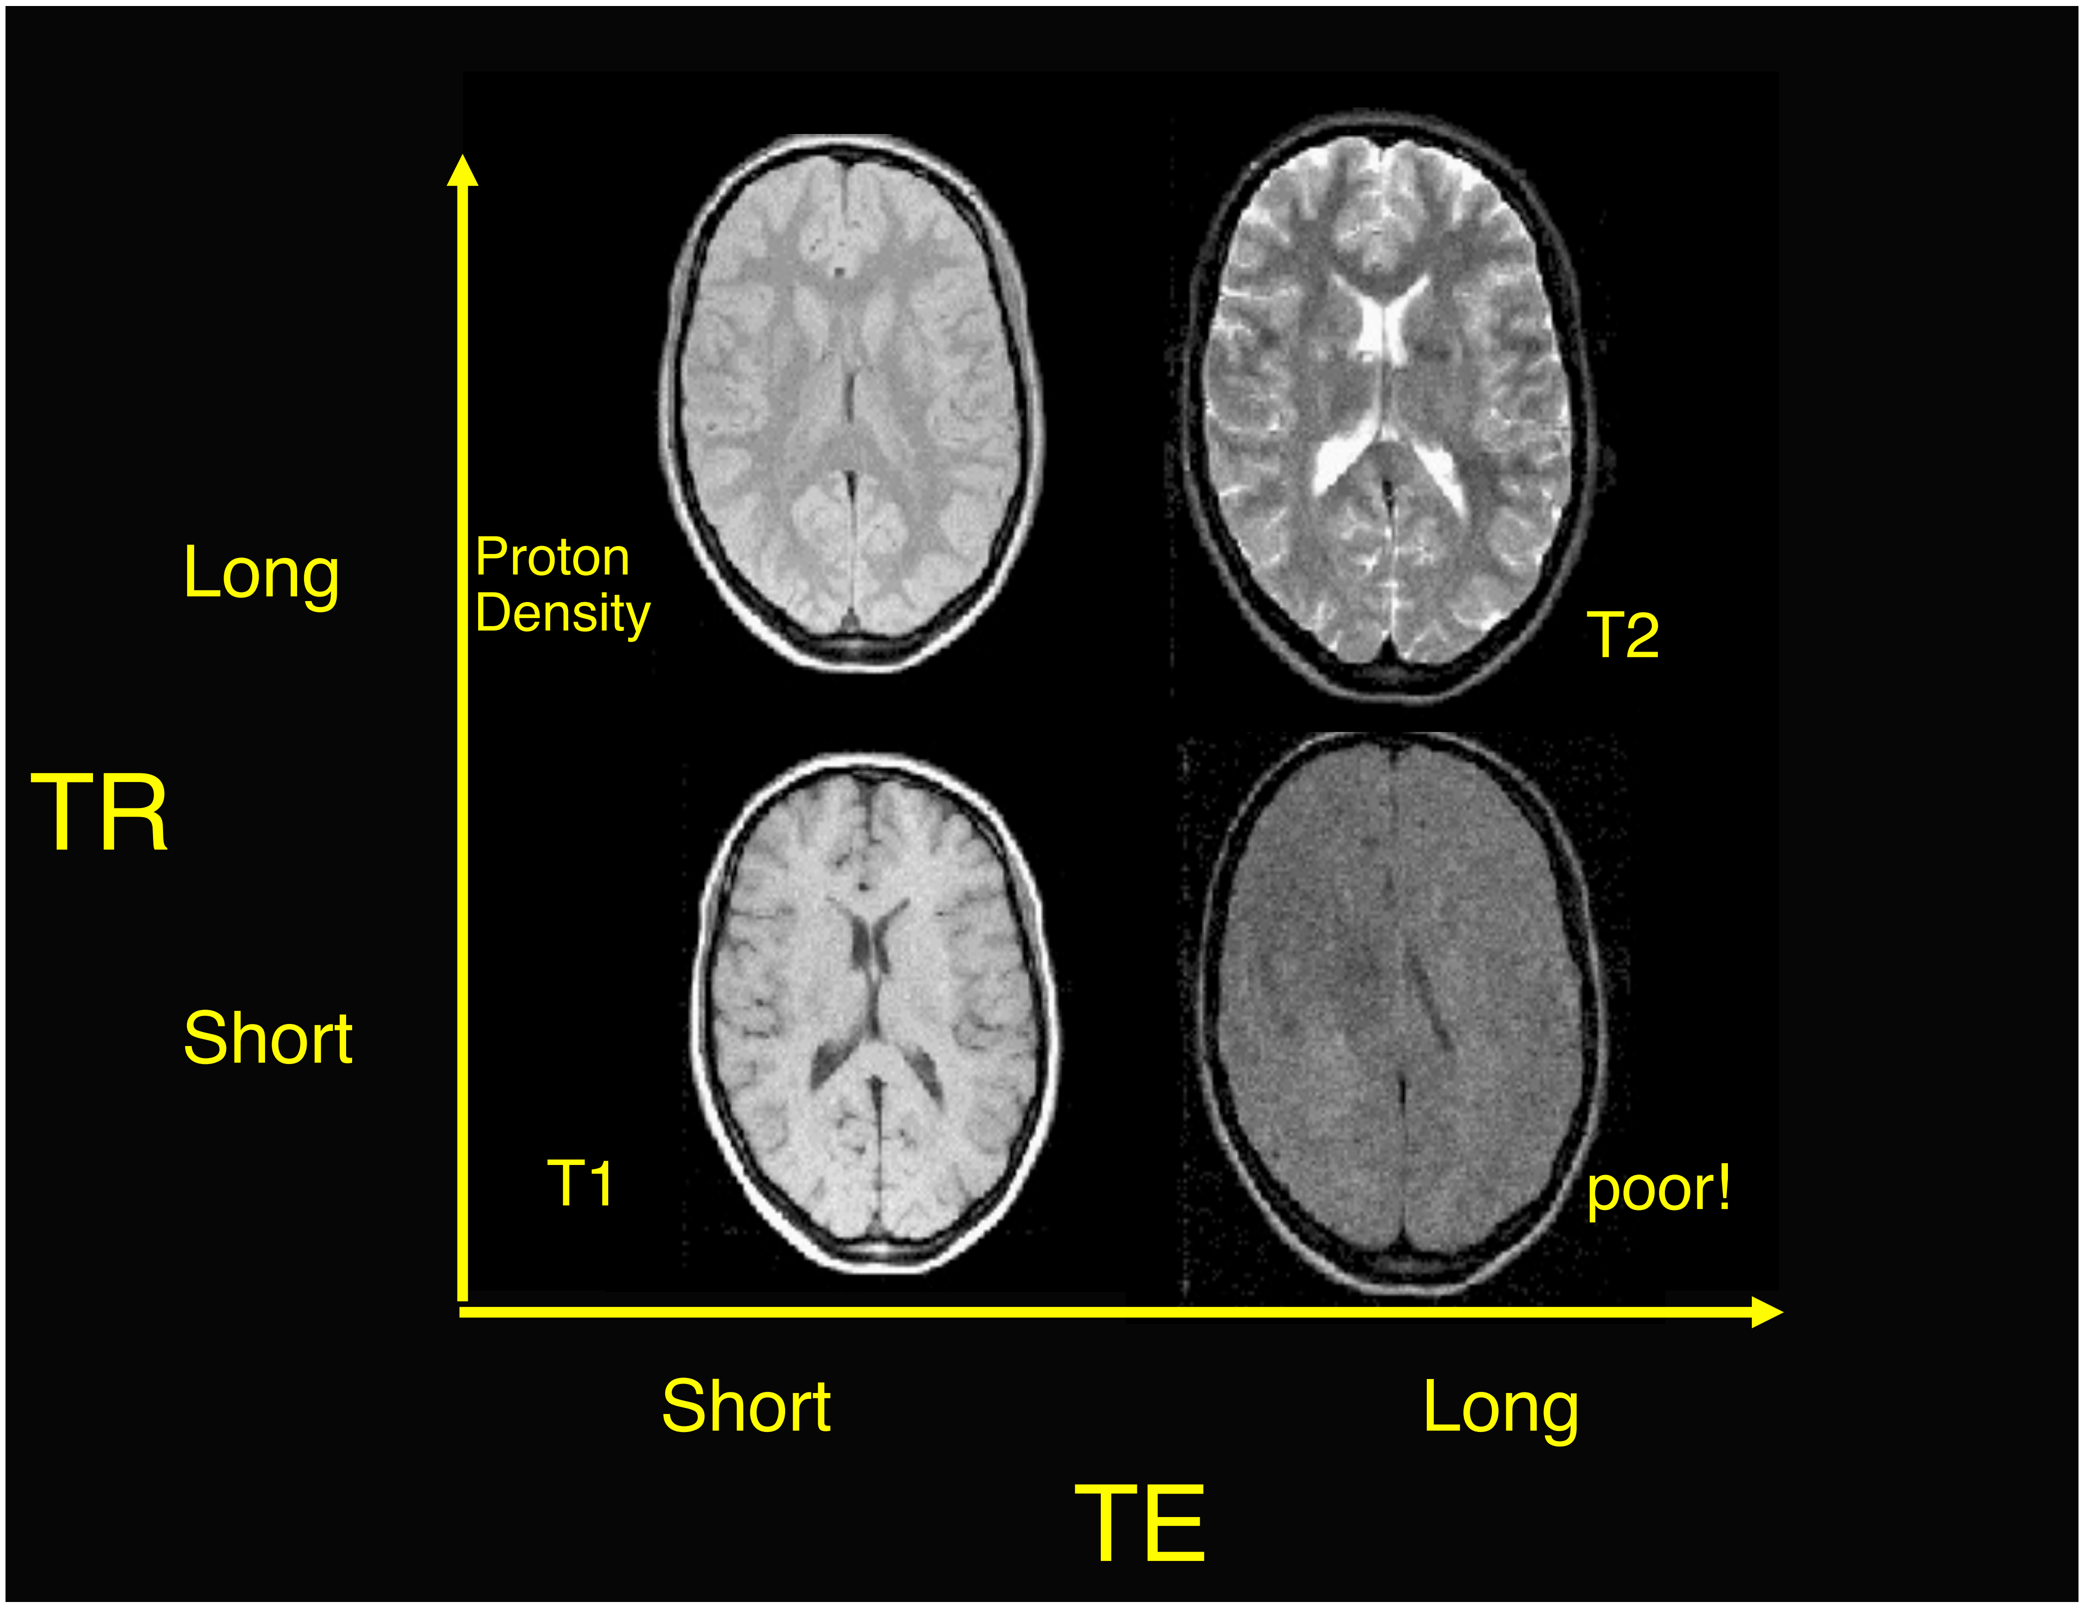

This is a summary or the imaging parameters TR and TE and the contrasts they bring out in the image.

Image encoding: gradients

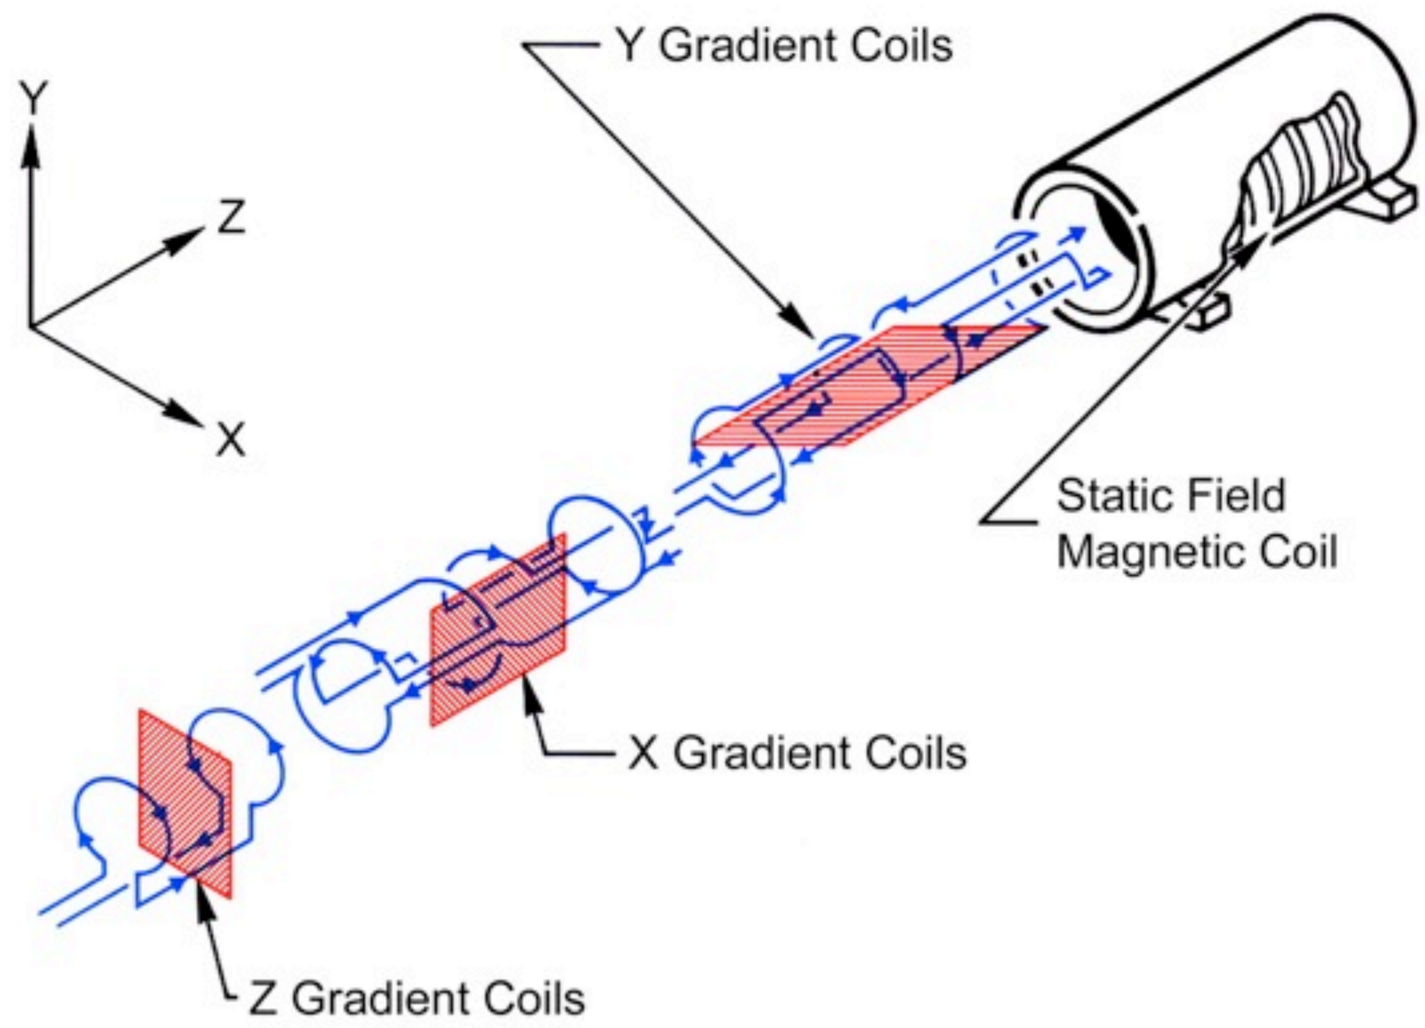

Gradient field add a linear space dependent B field to the original B_0 (still in z direction!)

Up until this point, I've really only shown you how to generate and detect a "signal": a time-varying voltage observed at the receive coil. However, we still don't know how to actually translate that signal into an image. For this, we need to go back to the scanner diagram and fill in an additional set of hardware: the gradient coils. A modern MR scanner has 3 gradient coils corresponding to the three spatial directions x, y and z.

They add a small B field (still in the z direction) whose magnitude depends on the position in a given spatial direction, e.g. x (or y or z).

Effect of a gradient field

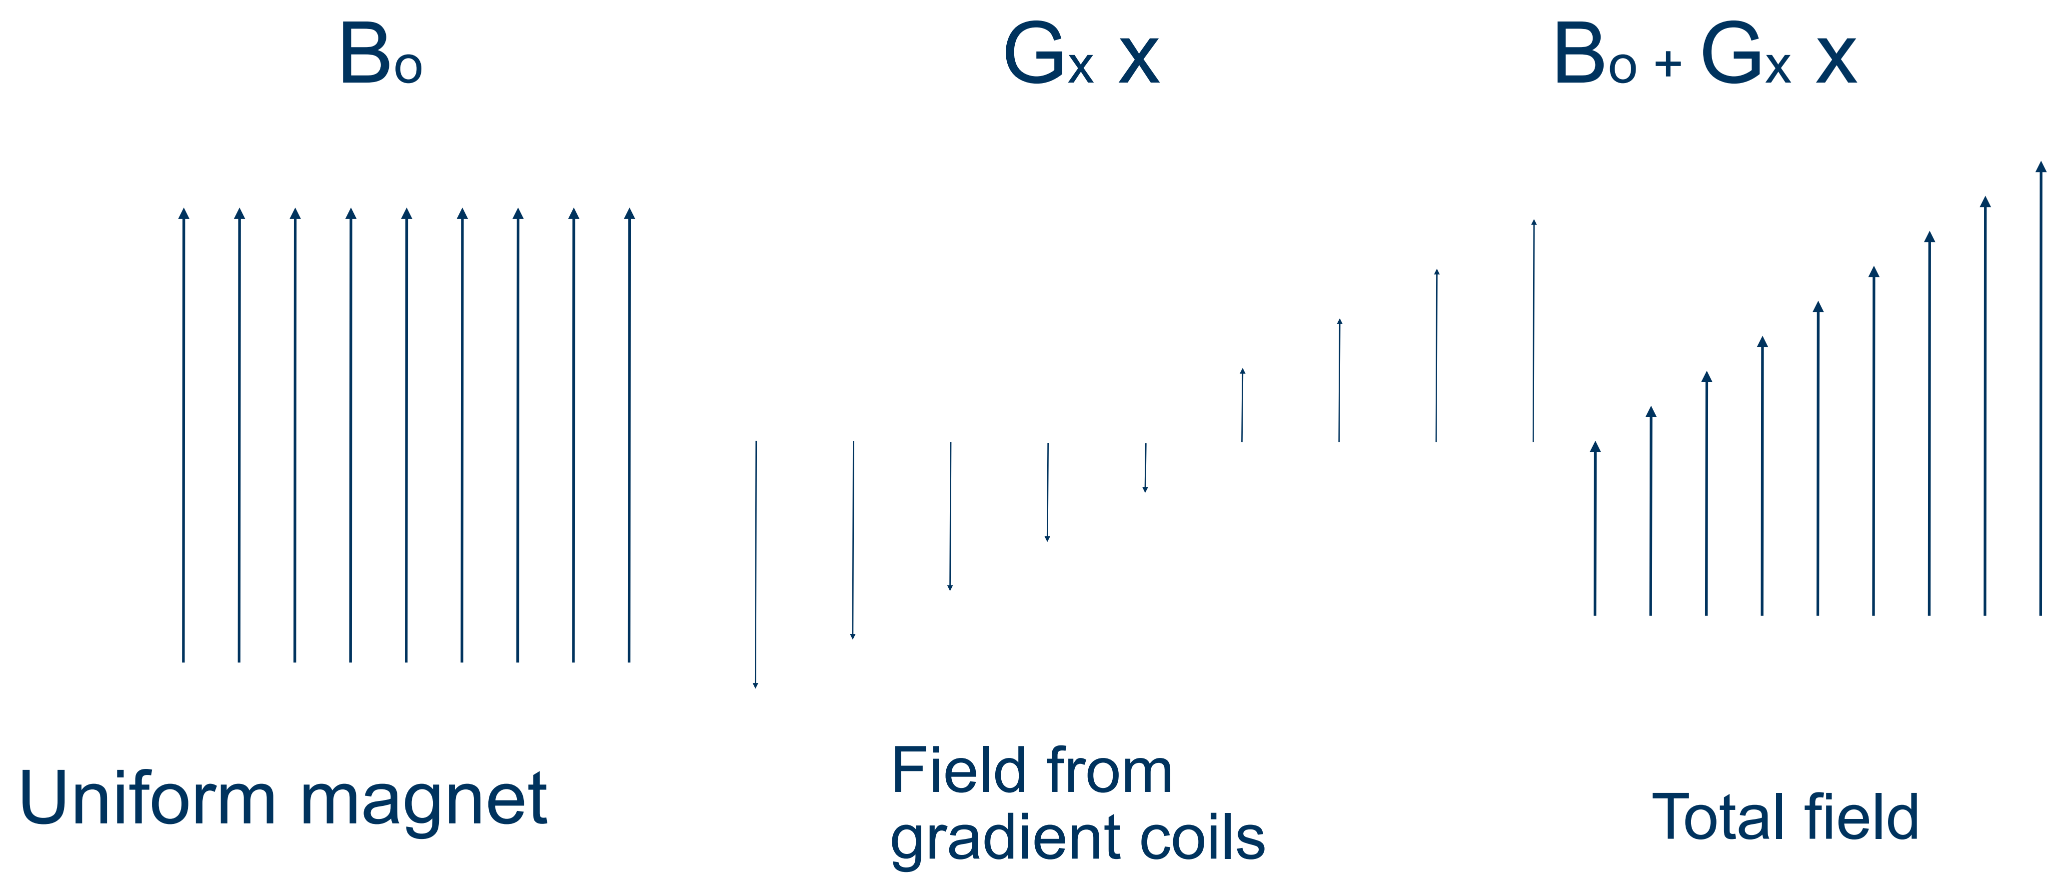

Let's look at the effect of an X gradient field. On the left we have the original B0 field, essentially spatially uniform.

In the middle, I show the additional field (scalled waaaay up in magnitude) contributed by the X gradient. Note that the additional field depends on where we are in X: large and negative at small X, then zero in the center of the magnet, then large and positive at large X (again.... all of this is still pointed along -- or against -- the Z direction).

On the right, we then see the vectorial sum of the two fields.

Frequency of precession w/ gradient

Precession frequency is proportional to the total B field (B_0 + gradient field)

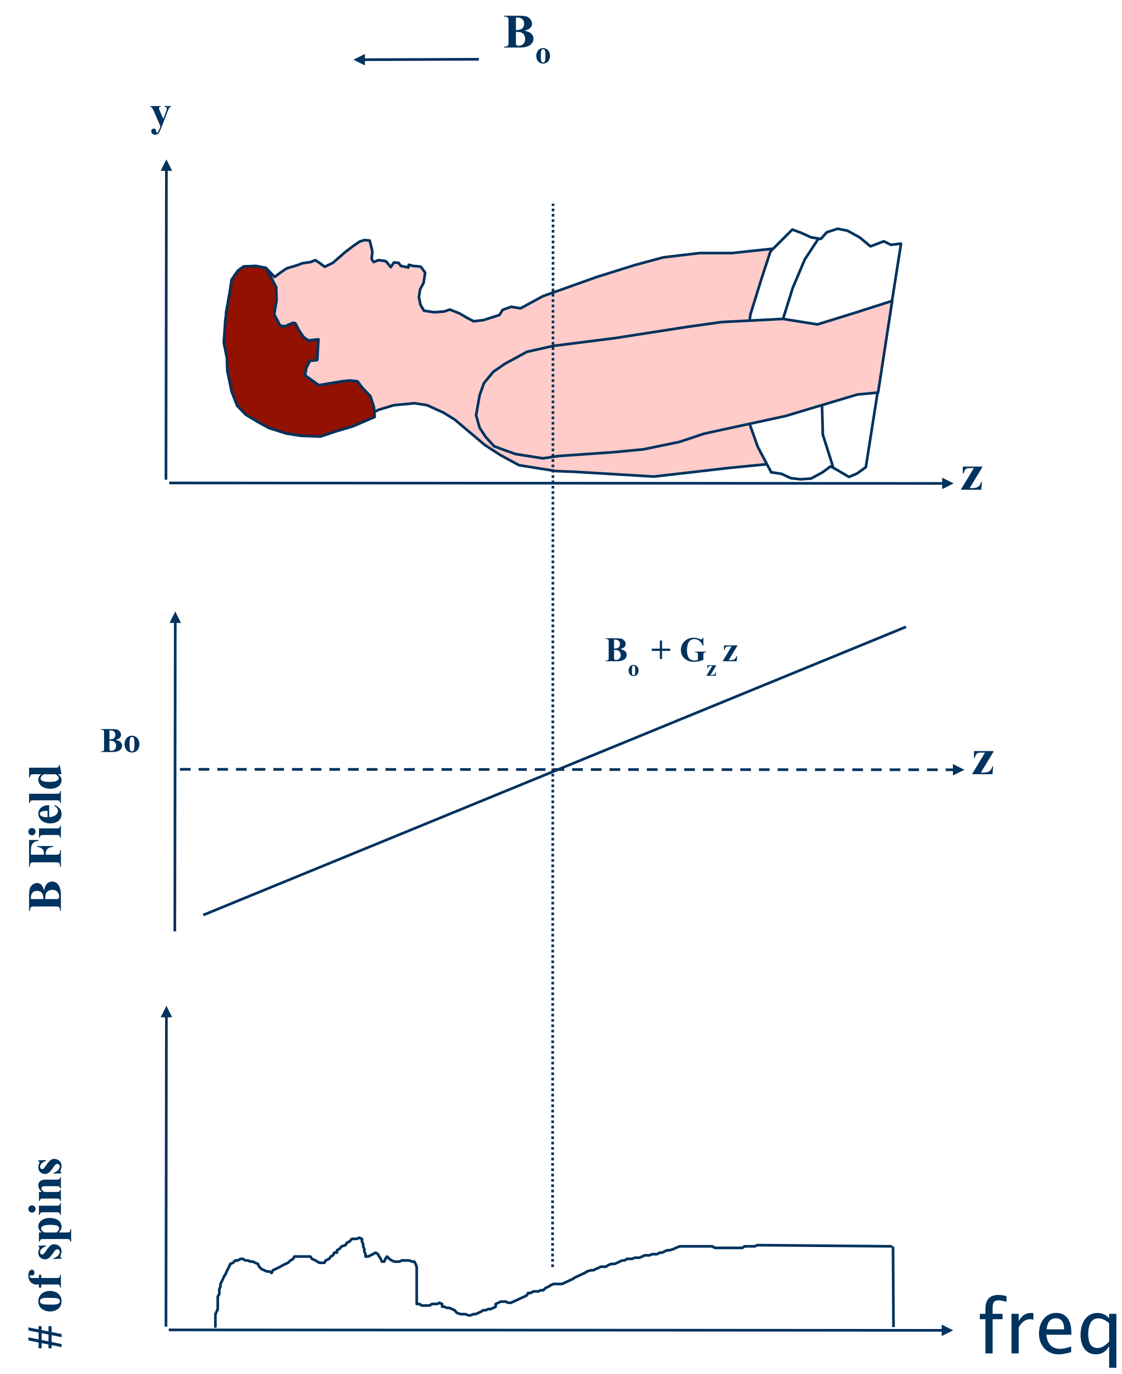

Now, recall that the frequency of precession of a proton in the MR magnet is dependent on the strengh.

Consider a Z gradient being turned on as the protons inside a subject are precessing. If we plot the number of spins precessing at a particular frequency, we see that we get what looks like a profile of our subject! We've created a correspondance between the position along Z and the frequency of precession. Already, we're getting close to imaging.

Slice selection (here w/ z gradient)

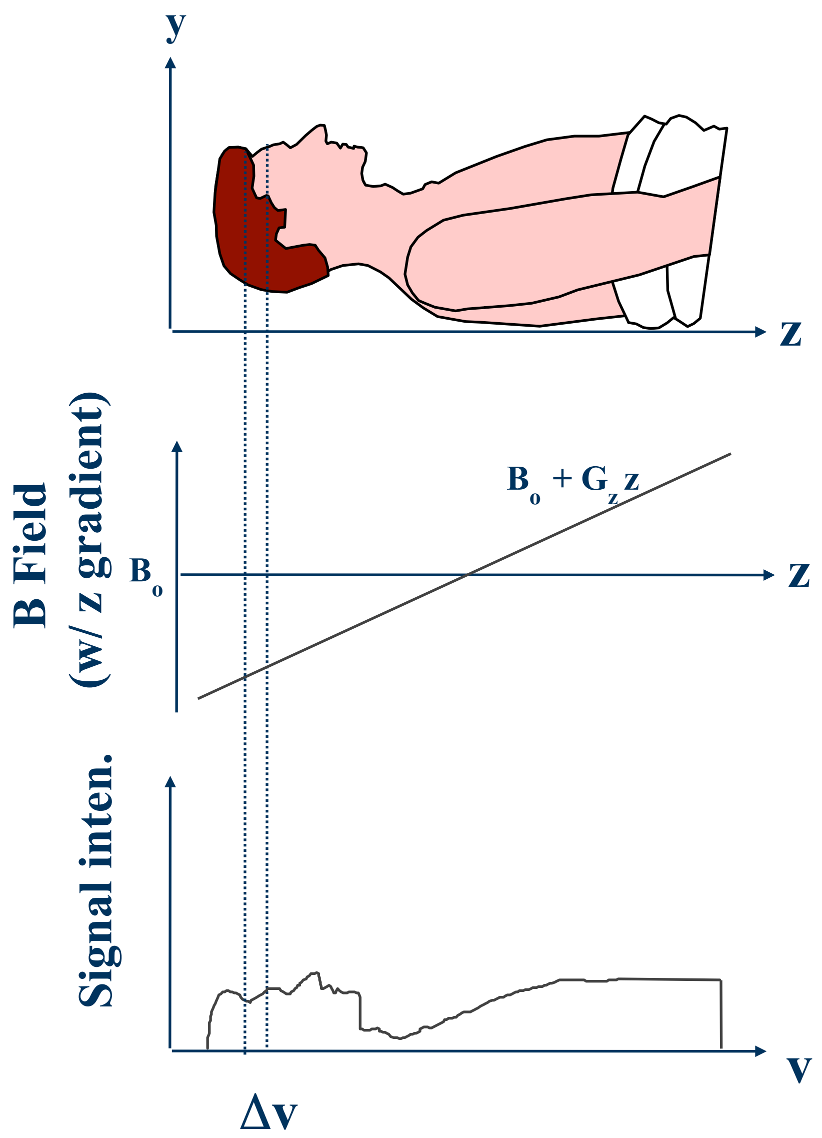

While the gradient is on, excite only band of frequencies.

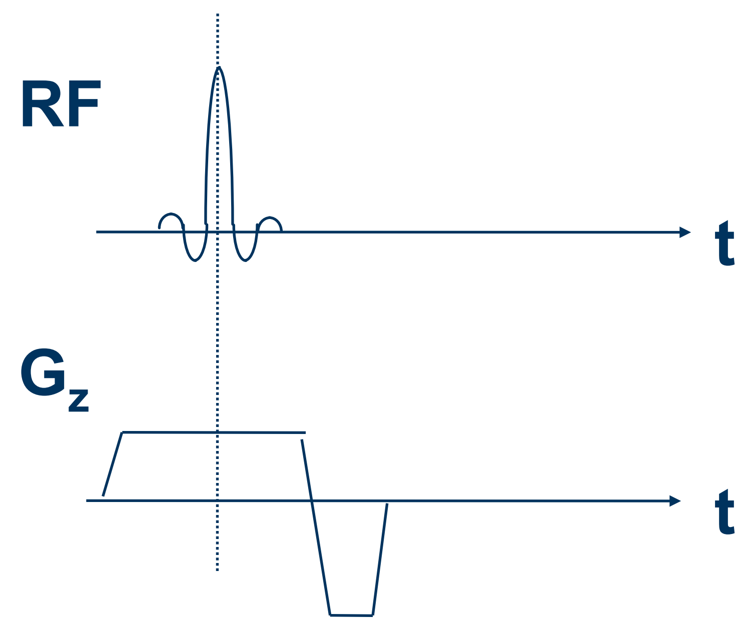

Now let's imagine turning on that Z gradient AT THE SAME TIME AS THE EXCITATION RF pulse. Let's further say that the RF pulse itself is not just a B1 field oscilating at ONE frequency, but a pulse that is shaped so as to excite all protons that are spinning WITHIN A RANGE OF FREQUENCIES. Now we have slice selection! Only the protons that are oscillating between some minimum and maximum frequencies will "feel" this excitation; and all those protons are located between some mimimum and maximum Z positions.

Slice selection has narrowed down the problem: we now detect a signal only from protons in a narrow Z slice, and we can repeat this to image as many slices as we like within time constraints. The next step will be to make it possible to identify which part of the signal comes from which positions in X and in Y.

Frequency encoding

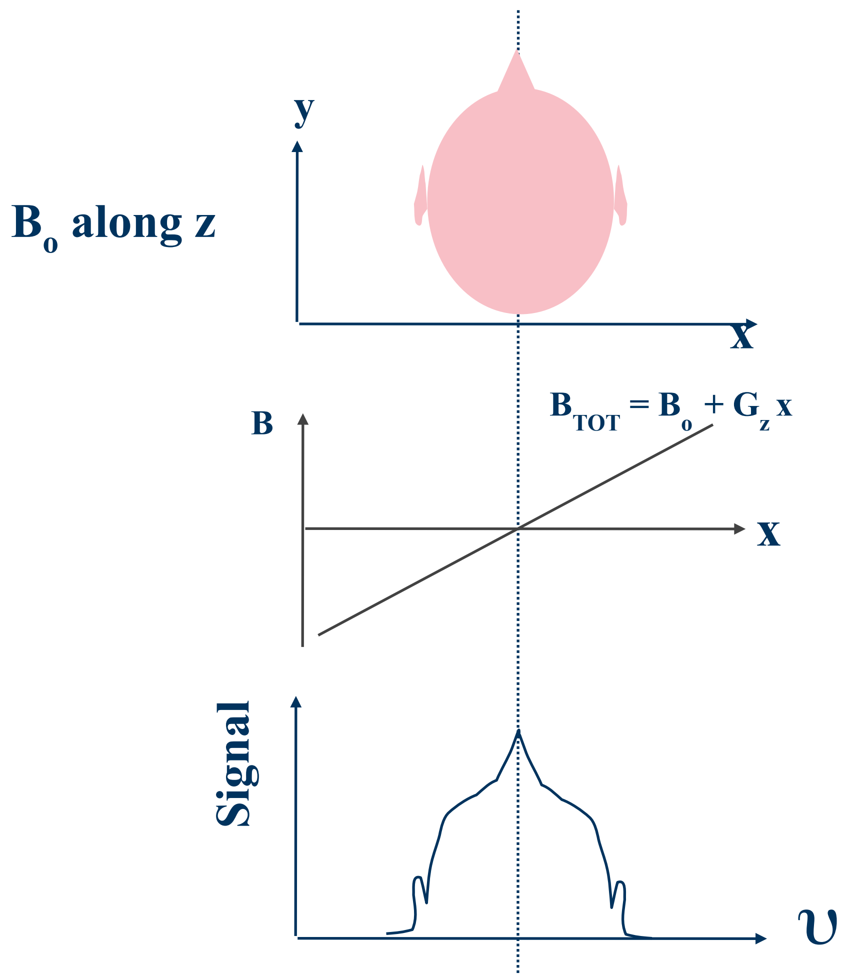

with gradient

"Frequency encoding"

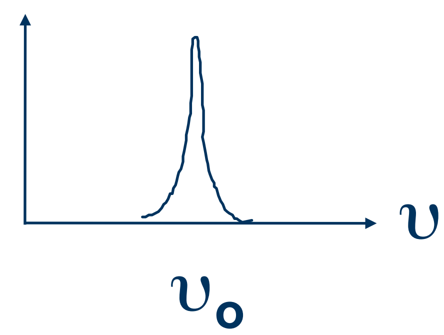

without gradient

We'll now use an X gradient during the ENCODING portion of our imaging sequence to obtain positional information along the X axis.

We've already excited a slice in Z, which I'm showing here from the top. Now, after the excitation, and during the encoding, we turn on a gradient in the X direction. Protons at different positions in X spin at different frequencies according to the gradient we've chosen to impose. Now the signal is a sum of oscillating curves, each one with a different frequency. At this point, we can "disentangle" the various contributions from different X positions, and reconstruct a 1D image that tells us where each contributing frequency is coming from.

Pulse sequence at this point

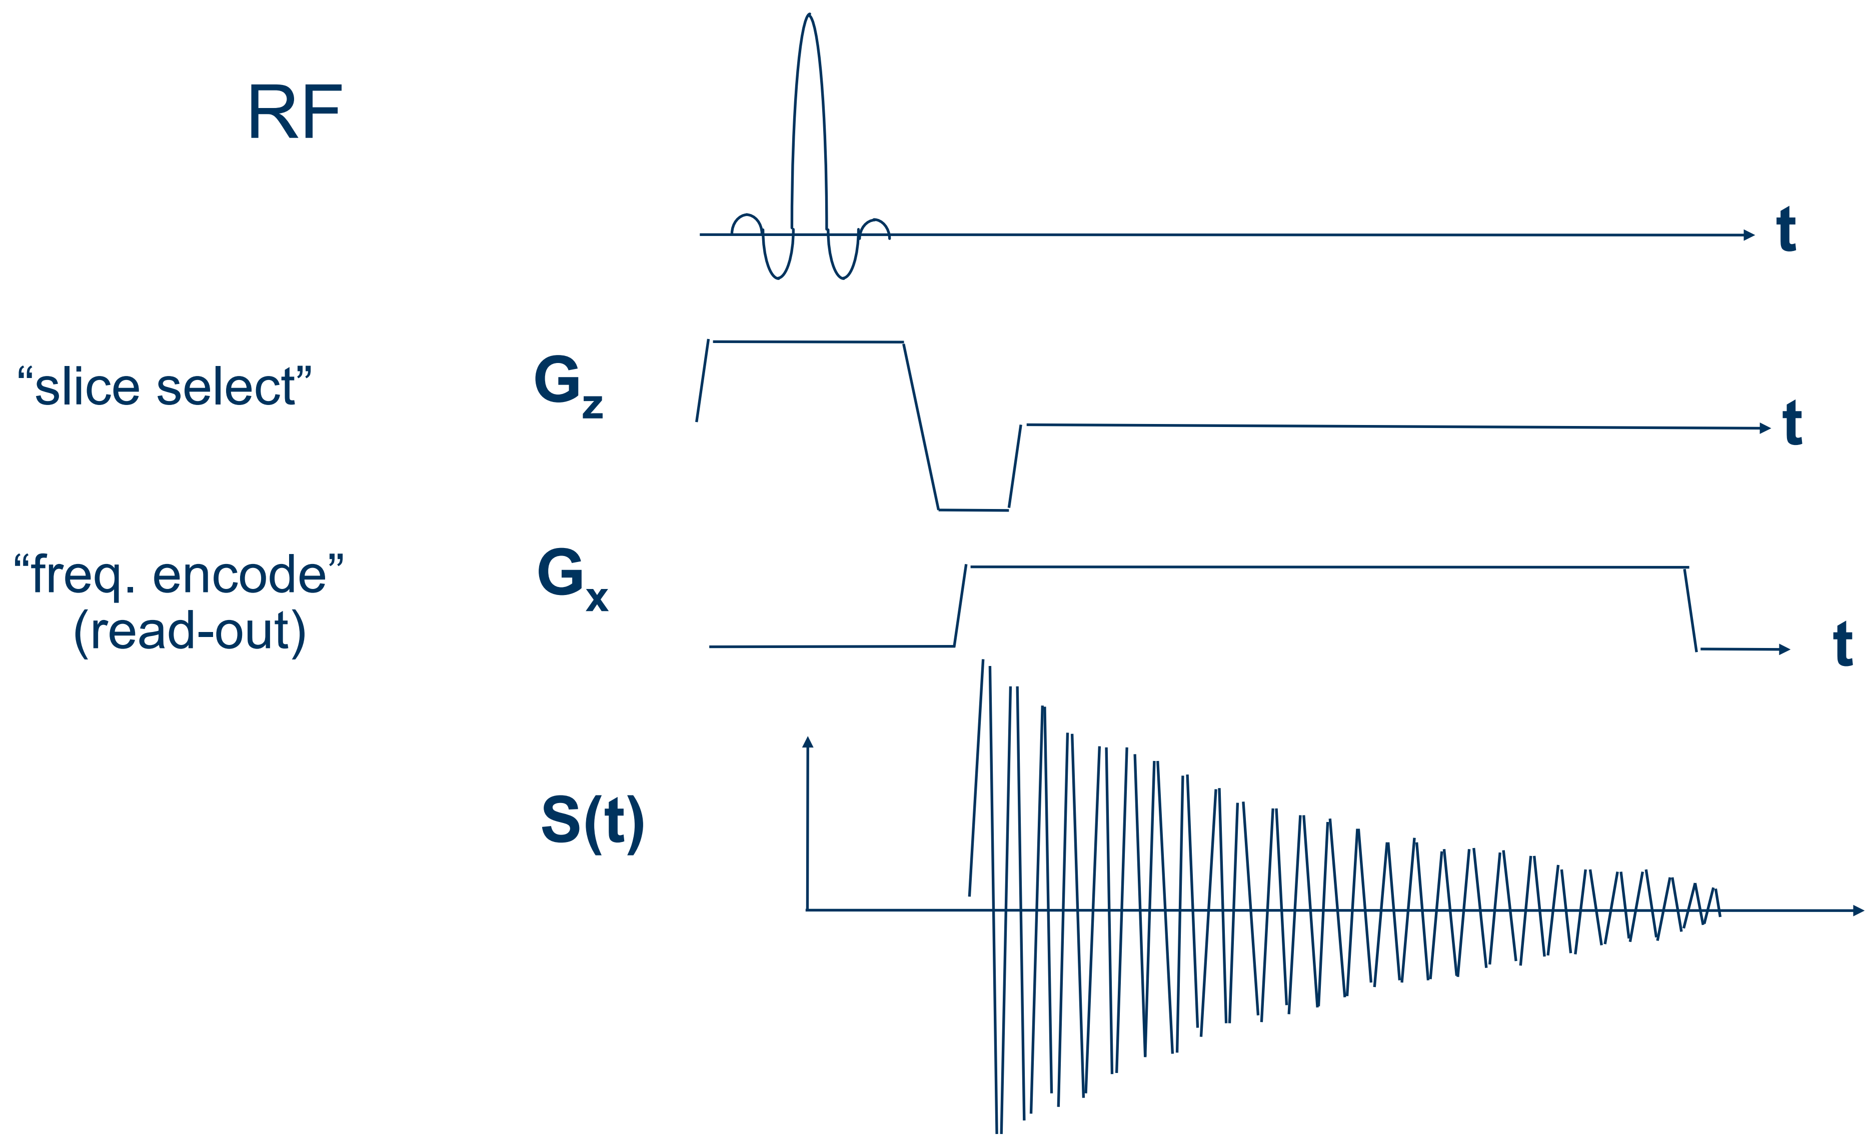

Let's look at the pulse sequence at this point, along with the signal we generate.

We start with an RF excitation, chosen to excite protons precessing with a frequency withing a particular range, itself chosen to correspond to the Z thickness of the slice we want to image. During that time, a Z gradient is turned on to enable this slice-selection to occur.

We then turn on an X gradient while we collect the signal, which encodes into that signal positional information along the X axis.

All that's missing is the Y axis...

Phase encode

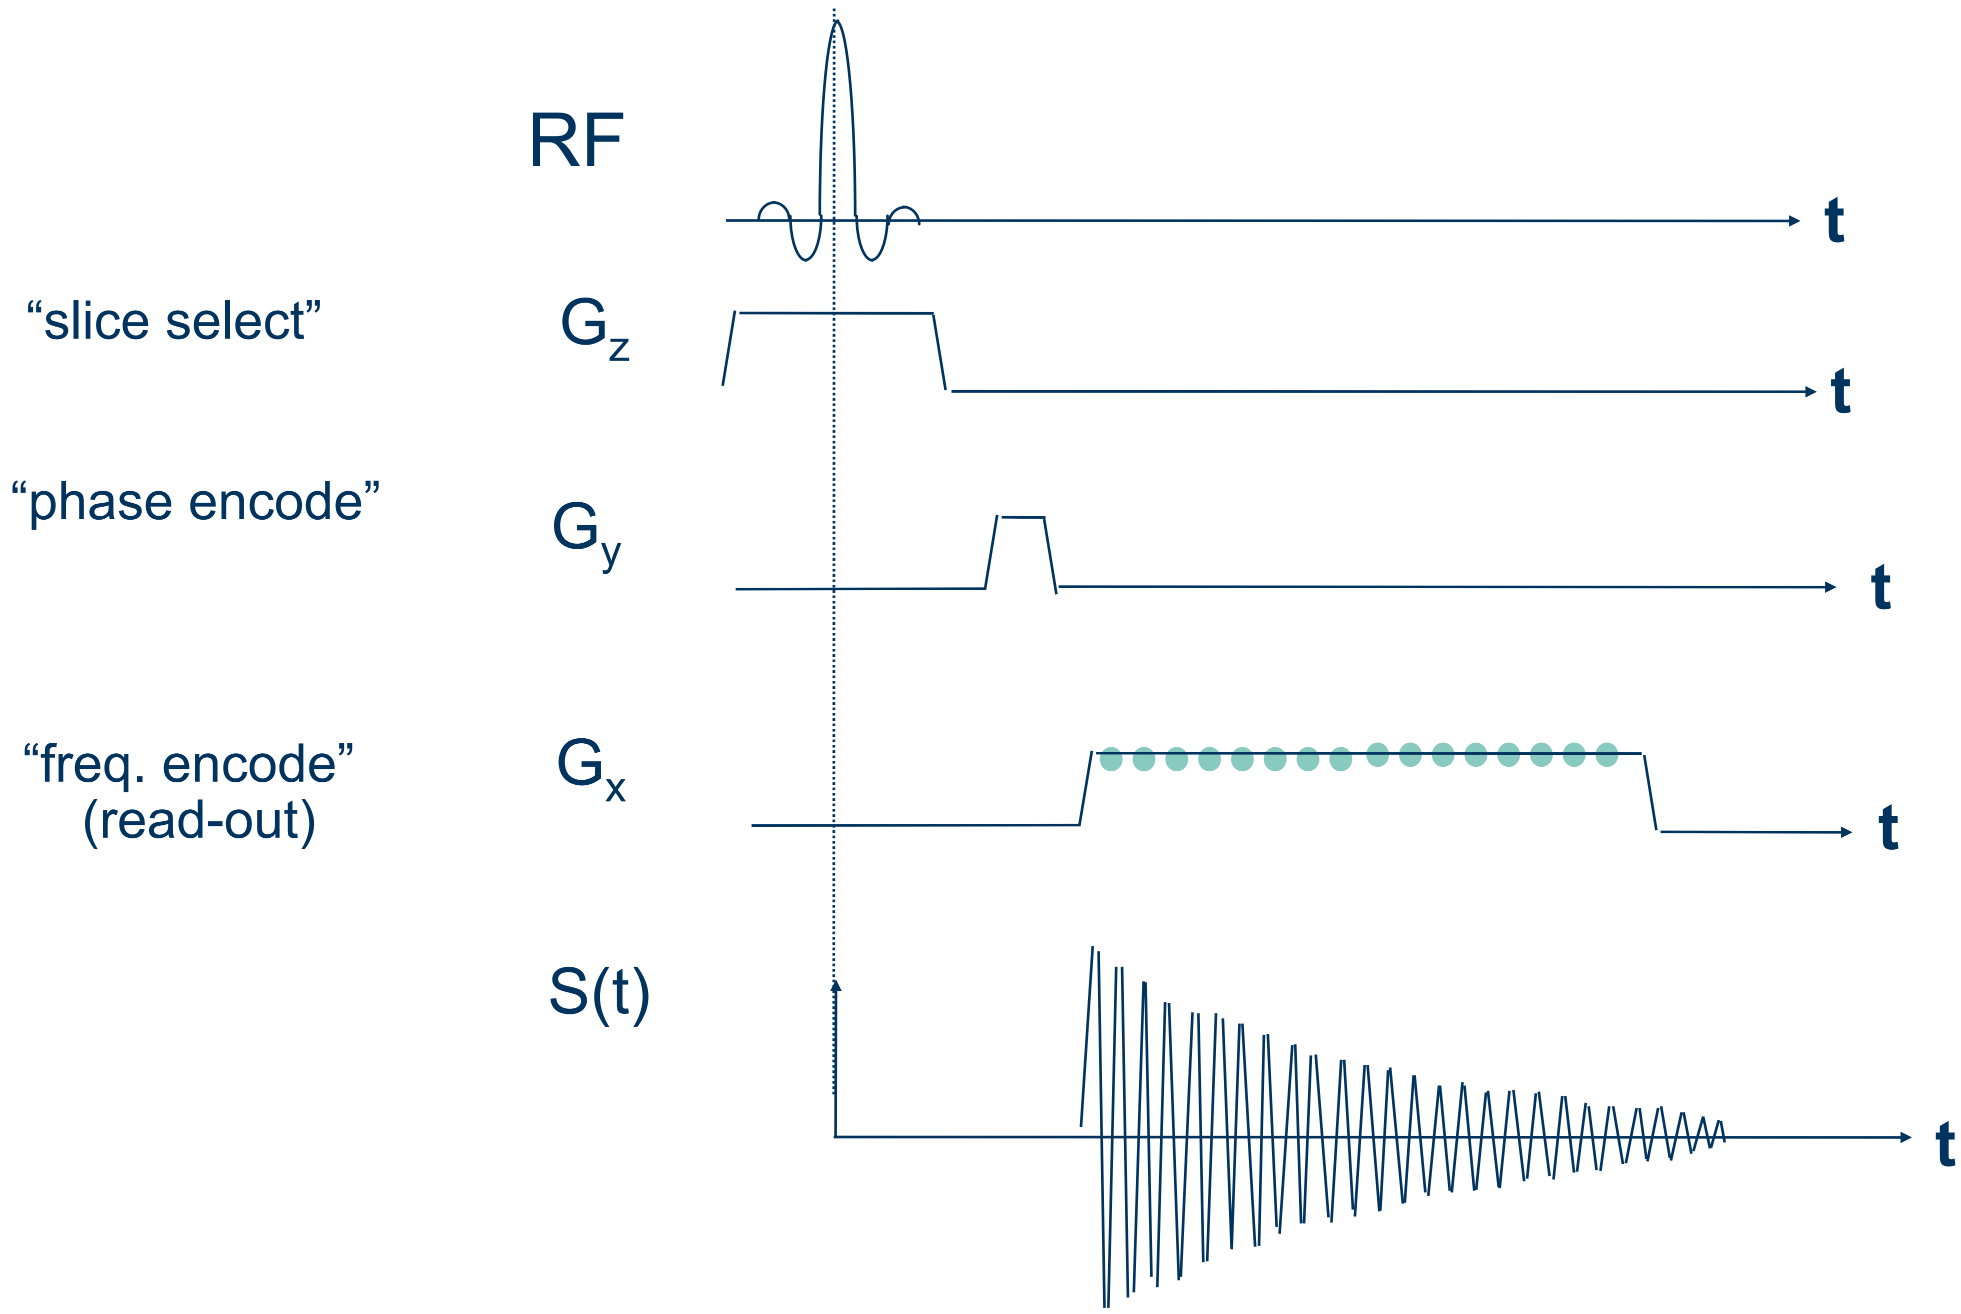

Phase encoding is what allows us to encode a 2D image, but unlike frequency encoding, is longer and more delicate to explain. It will therefore be outside the scope of this talk. Nevertheless, a quick mention is in order.

Here we are looking at a pulse sequence diagram very similar to the one we just looked at for a 1D frequency encoding. The only addition is the Y gradient that we briefly turn on after the excitation and before the encoding. What effect will that gradient have? It's not on during excitation, and therefore does not contribute to slice selection. It's on for a short time during which spins at different Y positions will precess with different frequencies, but then it's turned off before image encoding, by which time spins at different Y positions (same X) go back to precessing at the same frequency.

However, because the Y gradient is on for some nonzero time before we encode the image, spins at different Y positions with either fall behind or get ahead. Even when they go back to precessing at the same frequency as one another, this difference remains: they've now acquired a non-zero phase. This is what will allow us to now obtain Y positional information.

Phase encode....

For the full explanation, we need to look at this not in the image space, but in the spatial frequency space, i.e. k-space

Thanks!

Questions?