A Case Study in Small-Signal Measurement:

Detecting Faraday Rotation via Lock-In Detection

David A. Van Baak, TeachSpin, Inc. (version of July 2015)

0: Introduction

This laboratory exercise assumes that you know what Faraday rotation is, and that you have seen a qualitative demonstration that the phenomenon actually exists. (That is already quite an accomplishment, since this experiment provides direct experimental proof that a magnetic field in matter can affect the behavior of light – and this was, in 1845, the first-*ever* discovery of any relationship between light and magnetism.)

The purpose of this write-up is to show you how to measure Faraday rotation quantitatively, especially in the typical case in which the Faraday rotation is *small*. In such cases, the usual method of 'eyeball' demonstration fails to provide enough sensitivity. So this will show you how electronic detection capability can make your measurement of Faraday rotation much more sensitive than your original methods. In particular, it will introduce you to the power of lock-in detection for the selective measurement of weak signals that are polluted by electronic noise.

This treatment also pays some attention not just to the detection of weak signals amid noise, but also to their quantitative measurement. That is to say, it teaches you ways to get absolute, and not just relative, measurements out of a lock-in amplifier.

1. Faraday rotation: dc detection

The theory of Faraday rotation shows that a collimated beam of plane-polarized light, propagating in an isotropic medium parallel to a static and uniform magnetic field *B*, will undergo a rotation of its plane of polarization by angle

$$\theta = V\ B\ L\ .$$

Here *L* is the length of the sample, and *B* is the (assumed-uniform) value of the magnetic field. The material properties of the Faraday-rotator sample are captured in the Verdet constant *V*. In SI units, *V* has units of rad/(T·m). [You might also see it expressed in cgs units, typically in (minutes of arc)/(gauss·cm) or min/(Oe·cm).]

[See Section 5 of this write-up for the rich theory for the magnitude *V* is expected to have; there's a Becquerel theory which relates *V* to the dispersion of the Faraday-rotator material, derived from the function $n(\lambda)$, the dependence of the index of refraction *n* of the material on wavelength λ.]

Here's one way to measure *V* for a sample. For starters, we need plane-polarized light; we need a way to sense its (rotated) direction of polarization; we need a sample of material; and we need to immerse it in a static magnetic field, parallel to the light's direction of propagation. We achieve these by

- direct production of ≅650-nm light using a diode-laser source, in which the output light is born mostly polarized, and then fully polarized by an internal Polaroid;
- analysis of this light, downstream of the sample, by a rotatable sheet of Polaroid material, and the use of the 'extinction' technique to identify an angle just 90° different from the light's polarization;
- use of a fixed length *L* = 100 mm of a sample of a special material, a glass chosen for its comparatively large value of Verdet constant. In the TeachSpin experiment, the sample is a cylindrical rod of diameter 5 mm, made of Schott SF-57 glass (a glass optimized for high index of refraction, and high dispersion); and finally
- production of a magnetic field ***B*** to exist uniformly in the whole sample rod, and to have the correct (longitudinal) direction, by putting the sample inside the cylindrical bore of a simple solenoid.

With the solenoid built to be somewhat longer than the sample rod, we are to a good approximation in the infinite-solenoid limit, so ***B***'s magnitude takes on the textbook value

$$B = \mu_0\, n\, i\ ,$$

where $\mu_0 = 4\pi\times10^{-7}$ T·m/A and *n* (here) is the 'turn density', the number of turns of wire wound per unit length of the solenoid frame. In the TeachSpin solenoid, 1400 turns are wound (in 10 layers) into 6″ of length, so $n = 1400/0.1524\ \text{m} \cong 9168\ /\text{m}$, and we get a solenoid constant

$$k = \frac{B}{i} = \mu_0\, n \;\cong 0.0115\,\frac{T}{A} = 11.5\,\frac{mT}{A}.$$

Because of the finite length of the solenoid, the effective value of *k* (for the use of samples of length L = 100 mm) is closer to 11.1 mT/A. But notice that *B* will vary proportionally to current *i*; thus the predicted Faraday rotation ought also to be a linear function of the current *i*.

So here's the method:

1. We assume you've read the Manual, and know how to power the diode-laser source (using back-panel connections to the Power Audio Amplifier).

2. We assume you've followed instructions, and have aligned the laser, so that its beam is making a 'straight shot' along the center of the thin glass rod.

3. Now you're ready to install the Polaroid in its ring mount, and to search among 360° of rotation for (either of two) locations at which the intensity of the beam emerging beyond the Polaroid drops to very nearly *zero*. This condition, called 'extinction', is best seen by letting the beam fall onto a white card, and viewing that beam-spot on the card under subdued or suppressed room light.

4. When you've found an extinction condition, you know the 'pass direction' of the Polaroid has been adjusted to be exactly 90° away from the polarization direction of the light. As this 90° is a fixed offset, you need only record the angular reading of the Polaroid-mount's 'pointer' against the 0 - 360° angular scale.

5. With zero current in the solenoid, you expect zero Faraday rotation; but you should *not* expect a 0°-reading on the angular scale at the extinction condition. (The laser could have been installed into its mount with an arbitrary polarization direction; in fact, the Polaroid sheet you're using might have been installed at an arbitrary orientation relative to the pointer on its metal frame.) Despite such a non-zero angular offset, *changes* in the plane of the light's polarization will show up directly as changes in the angle-reading noted at extinction.

6. Once you've recorded the angle for extinction at $i = 0$, use a suitable power supply to send (say) a +3-A current through the solenoid. Find the new extinction-angle setting, and record its angular-scale reading.

7. Return to zero current, interchange the current leads, and now send a reversed, a -3-A current through the solenoid. Record the new angle setting needed for extinction. (Notice that it's inadvisable to break the current path while the current is running – the inductance of the solenoid will assuredly deliver a spark, which might damage the power supply.)

8. Suitable subtraction will give you $\Delta\theta$, the change in extinction-angle reading (which also equals the rotation of the plane of polarization) which accompanies a change in solenoid current.

Since your model is $\theta = V\,B\,L = V \cdot k\,i\,\cdot L$, you expect $\Delta\theta = V \cdot k\,\Delta i\,\cdot L$, so you can compute a Verdet constant

$$V = \frac{\Delta\theta}{k\,\Delta i\,L}\ .$$

If you measure (4 ± 1)° of rotation for a current of i = 3 A, you can infer $\Delta\theta$ = 0.07 rad, so

$$V \cong \frac{0.07\ rad}{\left(0.0111\frac{T}{A}\right)(3\ A)(0.100\ m)} = 21\frac{rad}{T\cdot m}\ .$$

The uncertainty in this value is totally dominated by the uncertainty in angle change $\Delta\theta$, which in turn is limited by the (± 1°?) readability of the 0-360° angular scale.

Note that if the Faraday rotation were even four-fold smaller than this, it would be hard to resolve *at all* by this ‘protractor’ method. What you’ll see in the rest of this exercise is that Faraday rotations a *1000*-fold smaller than those you’ve now seen are easy to detect and quantify by electronic methods.

2. Faraday rotation: large-signal electronic detection by oscilloscope methods

This section assumes that you've done Part 1, and have realized that the method of angle-scale readings you've used there lacks the resolution, and the sensitivity, that you'd like to have. In this section, you need only add a suitable photodetector, and a DMM and oscilloscope to monitor its output, to escape these limitations.

So in this Section you retain the laser source, the SF-57 glass sample, and the solenoid and its dc power supply; but you temporarily **remove** the Polaroid from its ring mount, and you add the photodetector described in the Manual. Before you actually mount it, look into its hooded entry-port, and locate the circular photo-sensitive area of its photodiode. This is a one-pixel 'bucket detector', not a camera, and its single output is a measure of the total optical power falling into that sensitive circle.

The photodiode converts red light to electrical current, with proportionality $i_{PD} \cong (0.6\ \text{A/W})\cdot I$, where I is the 'irradiance' of the incident light, the optical power it conveys. For your use of the TeachSpin detector, we suggest the choice by toggle-switch setting of the $R = 1\ \text{k}\Omega$ load resistor as a path for this current. Across that resistor, and the output cable, there is then developed a potential difference

$$V_{out} = i_{PD}\, R \cong \left(0.6\,\frac{\text{A}}{W}\right) I \cdot 1\ \text{k}\Omega = \left(600\,\frac{V}{W}\right) I = 600\ \text{mV} \cdot \frac{I}{1\ \text{mW}}\ .$$

Hence an optical power of $I = 0.5$ mW ought to give $V_{out} \cong 300$ mV.

The output voltage is very nearly linear in optical power I, provided that this V_{out} stays under 300 mV. [The toggle-switch load-resistor settings of $R = 3\ \text{k}\Omega$ and $10\ \text{k}\Omega$ are therefore suitable only with beams of much smaller optical power.]

Now connect the photodiode's cable to a suitable voltage-measuring tool. We suggest an oscilloscope, set to 50 mV/div vertically, and to about 500 ms/div horizontally, operating in the 'scroll' or 'roll' mode, so as to draw an ongoing graph of a voltage proportional to the optical power falling onto the detector.

Mount the photodetector into your TeachSpin optical bench, and set the height of the detector's post mount so that your V_{out} is maximized. Now make some slight adjustments to the alignment knobs of your laser source, and further maximize the V_{out} reading – this corresponds to getting the whole of the laser beam to fall into the active area of the detector. You should be able to "fall off a plateau" in V_{out} readings as you disturb the laser's alignment from this optimum – this corresponds to mis-aiming the laser so that part of its power I misses the photodiode's sensitive area.

When you've optimized this alignment, you may record the V_{out} value you've gotten, and you may admire its stability in time. But now you need to re-install the Polaroid in its ring mount, and you will then find that V_{out} becomes smaller, and of course it also becomes a function of the angle-setting θ on the Polaroid's angular scale. Take data through a 360° range at some suitable spacing, and note

- there are two minima or 'extinctions', giving $V_{out} \cong 0$, in the 0 - 360° range;
- there are two maxima (with V_{out} *smaller* than what you got in the Polaroid's absence) located just 90° away from the extinction angles;
- that a fair model for $V_{out}(\theta)$ is

$$V_{out}(\theta) = 0 + V_{max}\, cos^2(\theta - \theta_0) = \frac{V_{max}}{2}\,[1 + \cos 2(\theta - \theta_0)]\,.$$

Here we've assumed that the detector's output voltage V_{out} has two minima both consistent with zero, and has two maxima, of value called V_{max}, at angles we'll call θ_0 and $\theta_0 + \pi$, and we have further assumed that the detector is responding linearly with optical power, and that the Law of Malus is true for the Polaroid. A plot of this data, overlaid with a best fit, can give you 'best values' for V_{max} and θ_0.

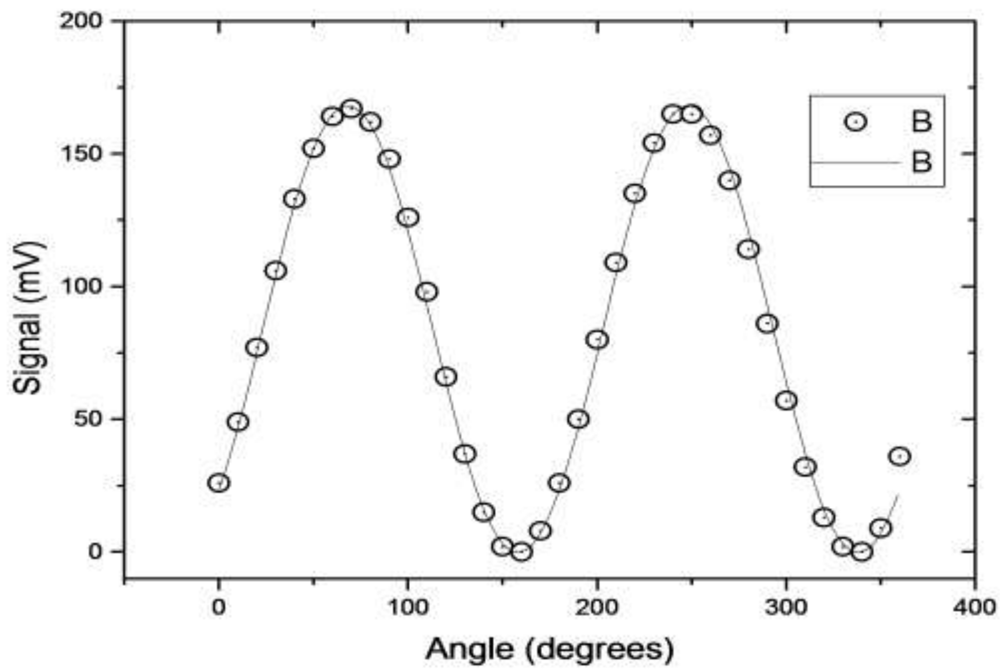

Your graph shows how V_{out} changes, for fixed angle-of-polarization of the light, as the <u>orientation of the Polaroid</u> is changed. But you may assume that this is also a valid model for the case of a fixed orientation of the Polaroid, when the <u>plane of polarization</u> of the light changes.

So looking at your graph, imagine that Faraday rotation is changing the angle θ of the light's polarization, relative to the angular scale on which you've found the angle θ_0. You can see that

- if you set the Polaroid to pass the maximal amount of light, any Faraday-induced changes in θ will slide you back-and-forth near the *maximum* in this plot, causing slight decreases from peak V_{out};
- if you set the Polaroid to the extinction angle, Faraday-induced changes in θ will slide you back-and-forth near the *minimum* of this plot, causing much more noticeable, but still slight, increases in V_{out} above 0;
- but if you set the Polaroid halfway between the 'pass' and 'extinction' conditions, you can operate on the points of maximum *slope* of your plot of V_{out} vs. θ, and get maximal sensitivity to Faraday rotation.

To see this algebraically, let ε be a small change in the light's angle of polarization (with ε measured in radians). Then you can show that the three choices mentioned above give

first, $V_{out}(\theta = \theta_0 + \epsilon) = \frac{V_{max}}{2}\,[1 + \cos 2\epsilon] \;\cong\; V_{max}[1 - 1\epsilon^2 + \ldots]$,

using a power-series expansion of the cosine function. Notice that the change of V_{out} from V_{max} is a decrease, and of *second* order in the small rotation angle ε.

second, $V_{out}\left(\theta = \theta_0 + \frac{\pi}{2} + \epsilon\right) = \frac{V_{max}}{2}\left[1 + \cos 2(\frac{\pi}{2} + \epsilon)\right] \;\cong\; V_{max}[0 + 1\epsilon^2 + \ldots]$,

again using a power-series expansion. The change are increases, again of second order in the rotation angle ε; they are more dramatic by eyeball detection only because you're comparing $0+1\varepsilon^2$ to 0 (rather than $1 - 1\varepsilon^2$ to 1).

but **third**, $V_{out}\left(\theta = \theta_0 + \frac{\pi}{4} + \epsilon\right) = \frac{V_{max}}{2}\left[1 + \cos 2(\frac{\pi}{4} + \epsilon)\right] \;\cong\; V_{max}\left[\frac{1}{2} - 1\epsilon^1 + \ldots\right]$,

using another power-series expansion. So in this case, you operate not at the curve's peaks (where $V_{out} \cong V_{max}$), nor at its valleys (where $V_{out} \cong 0$), but *halfway down the slopes* of the peaks, where $V_{out} \cong V_{max}/2$. The payoff is that you get ***first***-order sensitivity of V_{out} to changes in ε.

The distinction between first- and second-order dependence is crucial! Suppose V_{max} = 180 mV is the peak signal you get (at $\theta = \theta_0$ and at $\theta_0 + \pi$), and suppose that the Faraday effect causes a rotation of 0.57°, or ε = 0.01 rad.

Operating at peak, you get

$$V_{out} \rightarrow 180\, mV\, [1 - (0.01)^2] = 179.98\, mV\ ,$$

a decrease of 0.02 mV to be detected against a value of 180 mV.

Operating 'at valley' you get

$$V_{out} \rightarrow 180\, mV\, [0 + (0.01)^2] = 0.018\, mV\ ,$$

an increase of 0.02 mV to be detected against a value of near 0 mV.

But operating at points of maximum slope, you get

$$V_{out} \rightarrow 180\ mV\ [\tfrac{1}{2} - (0.01)^1] = 88.2\ mV\ ,$$

which is a change by -1.8 mV relative to a previous value of 90 mV.

So the changes in V_{out} you'd have to detect in the three cases are

a drop, by 0.02 mV, below 180 mV, or

a rise, by 0.02 mV, above a nominal value of zero, or

a change, by 1.80 mV, relative to 90 mV.

By a factor of *two orders of magnitude*, you win by operating 'on the slopes'.

Experimentally, this condition is very easy to arrange. You set the solenoid current to zero, and scan the Polaroid though 360°, making a good reading, to be called V_{max}, from among the various V_{out}-values you observe, and noting that V_{out} drops to near-zero at angles 90° away from these maxima. Then you merely go beyond one maximum, until your V_{out} drops just to *half* of V_{max}. You ought to find that you've gone 45° beyond a maximum.

Now your model for the signal comes from the 'third case' above:

$$V_{out}(\epsilon) = \tfrac{1}{2} V_{max} - V_{max} \cdot \epsilon\ ,$$

so that non-zero polarization rotations of size ε cause changes in V_{out} (relative to the steady value $V_{max}/2$) of

$$\Delta V_{out}(\epsilon) = -V_{max} \cdot \epsilon\ .$$

In practice, you'd use this to *measure* the rotation angles in radians indirectly, using

$$\epsilon = -\frac{\Delta V_{out}}{V_{max}}\ .$$

Now you're sensing Faraday rotation, but with the Polaroid left at a *fixed* location on its angular scale. Your resolution-in-angle is now limited <u>not</u> by the readability of the angular scale, but instead by the resolution of the electronics reading ΔV_{out}.

If V_{max} is 180 mV, and ε = 4°/(57.3°/rad) ≅ 0.07 rad, you expect V_{out} to vary, from a Faraday-free value of 90 mV, by the amount $V_{out} = - V_{max} \cdot \varepsilon \cong -12.6$ mV. If your DMM can read this even to ±0.1-mV resolution, you have a 1-part-in-126 voltage resolution limit, replacing your previous 1-degree-out-of-4 angular-resolution limit.

Notice that order-ε^1 detection capability also means that changing the *sign* of Faraday rotation changes the sign of ΔV_{out}. That is to say, if solenoid currents i = ±3 A give Faraday rotations of ±4° = ±0.07 rad, you expect changes in V_{out}, relative to 90 mV, of ±(-12.6) mV. Any effect that goes as order-ε^2 would *not* allow you to distinguish the effects of the sign of the solenoid current.

3. Faraday rotation: small-signal electronic detection by oscilloscope methods

This section assumes that you've set up a Faraday rotation apparatus, complete with laser source, an SF-57 glass sample, a Polaroid set 45° away from extinction, and a photodetector. It assumes that you have seen the changes ΔV_{out} in the photodetector signal that accompany changes Δi in the *dc* current you put into the solenoid. The novelty in this section is to motivate the substitution of a continually-changing, a sinusoidally-oscillating, *ac* current $i(t)$ for the former stepwise changes in dc current.

Recall that our model for the photodetector signal, under these conditions, is

$$V_{out} = \frac{1}{2}V_{max} - V_{max} \cdot \epsilon = \frac{1}{2}V_{max} - V_{max} \cdot V \cdot k\, i \cdot L \ .$$

Here V_{max} is the maximum photodetector signal you could get (easily checked by temporarily setting the Polaroid to be at the 'pass' direction, 90° away from extinction), and V and L are the Verdet constant and the length of the sample. The solenoid's coil constant is k, and i is the current you send through it.

You can best check this model's prediction that V_{out} is a linear function of i by recording V_{out} using a DMM (preferably one with 0.1-mV or better resolution) while changing i over a range of (say) -3 to +3 A. The slope of this graph is

$$\frac{\partial V_{out}}{\partial i} = -\, V_{max} \cdot V \cdot k \cdot L \ ,$$

so finding the slope, and knowing the values of V_{max}, k and L, you can solve for the Verdet constant quantitatively.

This is not a bad method, but it can be further improved. It is limited by two facts:

- The changes ΔV_{out} due to Δi are rather small, and they have to be detected in the face of the rather larger constant value $V_{\text{max}}/2$. If you look at $V_{\text{out}}(t)$ as a function of time on an oscilloscope, you'll get a vivid picture of the smallness of ΔV_{out} relative to the size of the steady value $V_{\text{max}}/2$. You can use the 'zero offset' of your 'scope's vertical axis to cancel out the constant, and then you can increase the 'scope's vertical-axis sensitivity to see the small ΔV_{out}. But this method only goes so far.
- The changes ΔV_{out} due to Δi that you want to detect have to compete with any *drift* in the 'constant' term $V_{\text{max}}/2$. The laser's power output is nearly, but not perfectly, stable, so in practice, the term $V_{\text{max}}/2$ exhibits some fluctuations. In your previous use of the dc power supply to run the solenoid, you would have dealt with such drift or fluctuations by taking repeated and alternating readings of V_{out} under the conditions of $i = +i_{\text{max}}$ and $i = -i_{\text{max}}$. But you can't alternate between those two situations as quickly as you'd like.

Now you're about to solve the problems of the smallness of ΔV_{out} relative to a constant, <u>and</u> the issue of drift and noise in the 'constant', by changing from a dc, to an **ac**, method of exciting the solenoid.

The procedure requires the temporary use of some sort of audio waveform generator, capable of generating sine waveforms, of amplitude about 0.5 V, of frequencies from <1 Hz to about 1 kHz. We suggest sending such a waveform *both* to ch. 1 of a dual-trace oscilloscope, *and* to the (dc-coupled) input of the TeachSpin Power Audio Amplifier, the very PAA you're already using to power the diode-laser source in your Faraday-rotation experiment. Now use two wires to connect the output of the PAA to the solenoid of your Faraday-rotation apparatus, and start with the 'Input Attenuator' or gain knob of the PAA set to near 0 on its 0 – 1 scale.

If you like, put a DMM-as-ammeter in series with the PAA-to-solenoid connection. Start with a low generator frequency (well under 1 Hz), and start with the DMM serving as a dc ammeter. If you choose a suitable gain setting on the PAA, you should see an actual alternation of readings, as positive and negative currents alternately flow through the solenoid.

Meanwhile, use ch. 2 of your 'scope to view the V_{out} signal from your photodetector, and choose a slow (≅1 sec/div) sweep speed for the 'scope. Your 'scope should show the 'cause' on ch. 1 (the generator signal which changes the solenoid current), and the 'effect' on ch. 2 (the photodetector signal showing changes ΔV_{out} relative to a steady value of $V_{max}/2$). You may need to raise the gain on the PAA to near-maximum, ie. use the largest solenoid currents you can get, to see this effect.

Once you've seen this, raise the generator frequency to about 30 Hz, and reconfigure your DMM-as-ammeter to be an *ac* ammeter. It will display the rms measure, I_{rms}, of the ac current you're sending through the solenoid. You can vary the 'gain' on the PAA to vary the strength of this ac current. Now you can also increase the sweep speed of your 'scope, and arrange for it to trigger on the ch. 1 'cause' signal. Next, use the ch. 2 vertical dc-offset control, and raise the ch. 2 sensitivity as much as you can, while still keeping the $V_{out}(t)$ signal on scale. Can you see the sinusoidal changes in $V_{out}(t)$, and can you see they occur at the same frequency as the generator's 'cause' signal?

Here are the advantages of going to a frequency of order 30 Hz. For one, your *ac* ammeter ought to respond properly to such frequencies (it would probably *not* do so at frequencies < 1 Hz; but it is surely designed to work correctly at frequencies of 50 or 60 Hz.) Better still, you can now configure your 'scope's ch. 2 input for ***ac***-coupling. This will *block* the constant dc value of $V_{max}/2$, but it will *pass* the ac part of the signal, the very signal ΔV_{out}(t) that you want to see.

So now you can eliminate the use of dc-offset from your ch. 2 input, and you can raise the sensitivity until the $\Delta V_{out}(t)$ waveform fills your vertical scale. The entire signal that you're seeing is due to Faraday rotation: here are two real-time ways to be sure of that:

- use a paper card to block the laser beam, back at the entrance to the solenoid. You should see the ch. 2 sinusoidal waveform collapse to a flat line. (Why?)
- remove that paper card, get the ch. 2 'effect' waveform back, and now dial the 'input attenuator' or gain knob on your PAA down to zero. Your ac-ammeter reading of solenoid current should drop to zero; that means the absence of any cause for Faraday rotation. And your ch. 2 'effect' waveform should again collapse to a (noisier) flat line.

(Why noisier this time? Because you now have, as you formerly lacked, a laser beam hitting the photodetector, so any fluctuations in the $V_{max}/2$ 'constant' signal will be showing up in this mode.)

Once you've confirmed these cause-and-effect connections, you can measure Faraday rotation quantitatively using this oscilloscope method. What you read on the ac ammeter is the rms measure of the solenoid current, I_{rms} , and the <u>amplitude</u> of your ac waveform is $I_{rms}\sqrt{2}$, so the solenoid current can be written as

$$i(t) = I_{rms}\sqrt{2}\ \cdot \sin(2\pi\ f\ t)\ .$$

Then your model for the experimental system predicts a detector signal

$$V_{out}(t) = \tfrac{1}{2}V_{max} - \ V_{max} \cdot V \cdot k\ i(t)\ \cdot L\ \ .$$

which means that the ac part of $V_{out}(t)$ is described by

$$\Delta V_{out}(t) = -\,V_{max} \cdot V \cdot k\ I_{rms}\sqrt{2}\,\sin\,(2\pi\ f\ t)\ \cdot L\ \ .$$

In other words, the amplitude of your 'scope's ch. 2 ac signal is predicted to be

$$V_{max} \cdot V \cdot k\ I_{rms}\sqrt{2}\ \cdot L\ \ ,$$

so if you know V_{max}, k, I_{rms}, and L, you can extract the Verdet constant from this amplitude reading.

Notice that the ch. 2 amplitude is predicted to vary <u>linearly</u> with I_{rms}, and you can check this claim. You need merely vary the gain setting on your PAA to get a variety of I_{rms}-values, and for each one, you can record the ch. 2 signal's amplitude. Do you see the cause and the effect change in proportion to each other? [You might also see another effect – at its highest gain settings, the PAA may generate a *distorted* version of a sine-wave current, and so you might see a distorted sine-wave Faraday-rotation signal – a tendency toward a flattening of the sine wave's maxima and minima is the clue you're looking for.]

[If you're tired of fixating on the Verdet constant, consider that downstream of your Polaroid analyzer, your laser beam has been 'amplitude modulated' by the signal being sent into your solenoid. What if you sent into the PAA's input, not a boring sinusoid from an audio generator, but an actual audio waveform from some music player? What would the ch. 2 'scope signal look like? How would you 'receive' this signal? How *far* could you convey your audio signal, optically?]

In these investigations, you may already have noticed that the ch. 2 'effect' waveform, for which you need to measure an amplitude, also shows electronic *noise* in addition to the waveform you're investigating. Here are three ways that you might improve this situation:

1. If you are triggering your 'scope properly, then your ch. 2 waveform is stable. But now that you can use your 'scope at full sensitivity on the vertical axis, it does make visible the accompanying random noise. If your 'scope has an <u>Average</u> function, this is

the time to try it out – you will get an average of N successive traces of the ch. 2 signal, which will partially average away the effects of noise, but leave the desired signal unchanged. [What are the advantages, and *dis*advantages, of making N larger?]

2. If your 'scope has a Measure function, you might be able to configure it to give a numerical reading for (say) the peak-to-peak excursion of your ch. 2 waveform. Half that peak-to-peak measure is the amplitude of $\Delta V_{out}(t)$ that you desire. The measure function might work best if you have ≅10 full cycles of oscillation displayed on your 'scope, and if you have applied the Average capability as well.

3. If your 'scope has an 'FFT' or fast-Fourier-transform function, then you can look at your ch. 2 signal in the frequency domain. To do this, remove the Average and Measure functions, and choose a horizontal sweep speed so that 20-50 full cycles of ch. 2 oscillation appear on-screen. Then the FFT function ought to give you the display of a spectrum of the input signal, which will reveal a spectral peak (at your generator frequency) atop background noise (distributed among all frequencies). The prominence of the spectral peak above the background is a good measure of how strong your Faraday signal really is. The vertical axis in FFT spectra is conventionally displayed on a logarithmic scale in dB(deciBels)/div, and a scale setting of 10 dB/div means each vertical division stands for ten times more spectral power (√10 times the amplitude) than the division below.

To confirm that your spectral peak is real, try

- blocking the laser beam – the peak (and some of the noise) should vanish.
- dialing down the gain of PAA – the peak should vanish, though the noise should remain.
- restoring the PAA to your old gain setting, and changing the generator frequency – the peak should move sideways, to occur at the new frequency you've selected.

While you're working in the frequency domain, find out how small a gain setting on the PAA (and how small a consequent I_{rms} value) you can choose, that will still result in a recognizable spectral peak corresponding to Faraday rotation.

Furthermore, in this frequency-domain view, you can also look for 'competition'. Because your detector is seeing some room light, and because that light is likely to be modulated at 100 or 120 Hz (why?), there could very well be another spectral peak in your FFT display. Or, if there is electrical (rather than optical) interference, there could be a peak at 50 or 60 Hz. The point is that you get to *choose* where in frequency space you put your signal peak, and thus you can choose the oscillator frequency to *avoid* such excess noise peaks. You can also avoid the 'excess low-frequency noise' which exists, and might be visible at frequencies of 20 Hz and lower.

One thing you probably can*not* do with a basic oscilloscope FFT display is to make a quantitative measurement of the amplitude corresponding to the spectral peak that you have seen. The typical oscilloscope can only measure this amplitude quantitatively in the *time domain*, and you have perhaps already measured it that way. But as you dial down the PAA gain and consequently lower the value of I_{rms} on your meter, you may have seen the Faraday-rotation signal apparently disappear under the noise in your time-domain view of the waveform. But an FFT view should persuade you the signal is nevertheless still present, and readily visible, in the frequency domain.

So the next section takes up a specialized task that often arises in physics – how can one accurately measure the amplitude of a sinusoidal signal of a predictable frequency and phase, even when it's buried under noise? The tool of choice for such a task is the lock-in amplifier, whose use in this instance of Faraday rotation we will describe in the next section. You'll find that while an oscilloscope can barely display the presence of the milliVolt-level sinusoid, a lock-in amplifier makes it easy to measure accurately the amplitude of even a *micro*Volt-level sinusoidal signal.

4. Faraday rotation: small-signal electronic detection by lock-in methods

This section assumes that you have completed Section 3 above, and now wish to replace your oscilloscope with a lock-in amplifier as the tool of choice for measuring the amplitude of a sinusoidal signal. In what follows, we describe the use of TeachSpin's education-optimized Signal-Processor/Lock-In amplifier, in part because this designed-for-teaching instrument allows you to *follow* the progression of your Faraday signal through its circuitry. But whether you use the TeachSpin or another lock-in, recall these facts:

- a lock-in amplifier's output is a dc voltage, which is proportional to the amplitude of the input sinusoidal waveform you are measuring;
- but to know *what* frequency (and phase) of sinusoid to quantify, the lock-in needs to be provided with a <u>second</u> input signal, called the reference input.

Typically the input signal presented to the lock-in is weak and/or noisy; it comes from the 'effect' waveform emerging from an experiment. Typically the reference signal is derived from a copy of the 'cause' waveform in the same experiment, so there's no reason for it to be either small or noisy.

We'll now describe how to set up the TeachSpin lock-in to accomplish this measurement of amplitude. The case you should try first is a Faraday-rotation experiment, in which you get a maximal Faraday-rotation signal, through the use of the SF-57 glass sample, and as big an amplitude of ac solenoid drive as you can achieve. We suggest this case because

> **we firmly believe that your first encounter with the lock-in technique ought to be in a case where you can at least *see* the signal (however much it's afflicted with noise) on a familiar 'scope display.**

Once you've learn how a lock-in works, you will be able to understand how and why it goes on working, even when the signals it's detecting and quantifying are *orders of magnitude* too weak to see on a 'scope.

First, we suggest you make one change in your Faraday rotation set-up. In place of, or in series with, your ac-ammeter monitoring the solenoid current, add a 1-Ω, ½-W resistor.

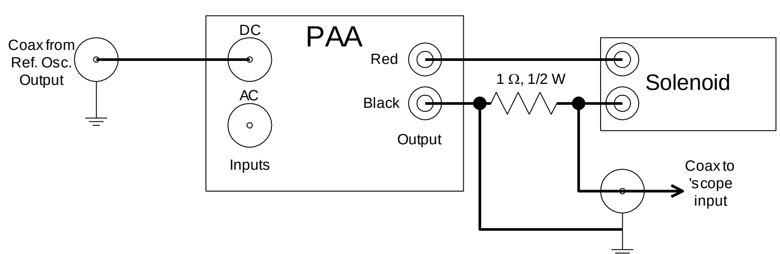

Connecting a current-monitoring resistor into the Power Audio Amp's drive of the solenoid. Note that the grounds enforced by the input, and the output, coaxial connections are *compatible* for the wiring shown.

The potential difference across this resistor will provide a voltage waveform $V_i(t) = (1\ \Omega) \cdot i(t)$, which serves as a *surrogate* for the actual current waveform in your solenoid. In fact, you should first use ch. 2 of your 'scope to look at this very waveform. Is it sinusoidal? or is it distorted or 'clipped' by PAA limitations when using the highest PAA gains? How big can you make this waveform? Is it in phase with the drive waveform you have displayed on ch. 1? (It might fail to be – after all, you are using the PAA to drive a load with both inductance and resistance, so phase shifts are to be expected.) In lock-in detection, phases like this *matter*, and this new $V_i(t)$ waveform will help you get them correct. In particular, the usefulness of the $V_i(t)$ waveform is that, as the immediate 'cause' of the Faraday rotation, it ought to have the same phase as your weak and noisy $\Delta V_{out}(t)$ waveform from Faraday rotation.

Start by turning on your TeachSpin lock-in, and find its Reference Oscillator section. Configure that to be a sine-wave oscillator at frequency about 30 Hz, and look at its (left-hand) variable-amplitude output. You'll now be using this waveform in place of your previous audio generator to drive the PAA and the solenoid. Set the toggle switch to the ×1 attenuator setting, and set the variable amplitude of this left-hand output to give about an amplitude of about 0.5 V, and use it to drive the PAA. Also send the Reference Oscillator's *right*-hand output to ch. 1 of your 'scope, and trigger the 'scope on this waveform.

Now use the 'scope's ch. 2 input, configured for ac-coupling, to look, alternately, at $V_i(t)$ [a surrogate for the 'cause' waveform] and $\Delta V_{out}(t)$ [the photodiode-generated 'effect' waveform]. These will have quite distinct amplitudes, with $V_i(t)$ large, and $\Delta V_{out}(t)$ small. They will also differ in 'noisiness'. But both will have the same frequency as the oscillator, and both will have a stable phase relative to the oscillator.

To see that 'cause' and 'effect' are in phase with *each other*, now send your Oscillator's right-hand output waveform not to ch. 1, but to the External Trigger input of your 'scope. Arrange for the 'scope to trigger on this external input, so the Oscillator waveform will be unseen but still in control of the 'scope's triggering. Now use ch. 1 to view $V_i(t)$, and ch. 2 to view $\Delta V_{out}(t)$, and set the horizontal sweep speed so that 2-3 cycles of the these waveforms fill the screen. Can you see that they're in phase? This is a very direct time-domain view of the 'cause' and 'effect' in your experiment. Again, you could use an Average function on your 'scope to lessen the effects of noise.

[Actually, it's a *coin toss* whether in fact your two waveforms come out at 0°, rather than 180°, phase difference. If you get, but dislike, the 180° outcome, rotate your Polaroid to 45° on the <u>other</u> side of a maximum of $V_{out}(\theta)$, so as to operate on the *opposite* slope, ie. on the *other* side-slope, of the $V_{out}(\theta)$ curve you've previously graphed. Perhaps you can intuit why this should make a sign difference in your signal. (Rotating that polarizer will also allow you to see that $\Delta V_{out}(t)$ vanishes when you operate at either the peaks, or the valleys, of the $V_{out}(\theta)$ curve.)

Alternatively, if you prefer not to change your Polaroid setting, you can cure a 180° phase shift merely by interchanging, <u>at the binding posts of the solenoid</u>, the two electrical connections you've made.]

[One more fine point: to get a correct view of the phase relationship of ch. 1 and ch. 2 signals, you'll need to configure your 'scope so that its two input channels are either both dc-coupled, or both ac-coupled. Otherwise, there'd be a relative phase shift caused by the dc-blocking high-pass-filter that's present in one channel but absent in the other.]

Now that you've seen these waveforms on the 'scope, it's time to use the lock-in to measure, in succession, the amplitudes of the waveforms $V_i(t)$ and $\Delta V_{out}(t)$. Recall that the current-monitoring signal $V_i(t)$ has amplitude

$A_i = (1\ \Omega) \cdot I_0$ where I_0 is the amplitude of the solenoid-current waveform,

while the optical signal emerging from your photodetector has an ac component of amplitude

$$A_{out} = V_{max} \cdot V \cdot k\, I_0 \cdot L\ .$$

So when you have measured these two amplitudes using the lock-in (as you might already have, using the 'scope), you will be able to extract the Verdet constant from them as

$$V = \frac{A_{out}}{V_{max} \cdot k\, I_0 \cdot L} = \frac{A_{out}}{V_{max} \cdot k\ (A_i / 1\ \Omega) \cdot L} = \frac{A_{out}}{A_i} \cdot \frac{1\ \Omega}{V_{max} \cdot k\ L}\ .$$

Now we describe is the process necessary for measuring amplitudes using the lock-in; we suggest you first try this on $\Delta V_{out}(t)$, the waveform that suffers most from being rather small, and somewhat noisy. For best pedagogical results, we suggest that you do the steps below not as a blind or 'paint-by-number' exercise, but as a case of active learning, in which you use a 'scope to verify what is happening at each step, and in which you also get to think about why it's happening, and how this might be useful.

1. Find the Pre-Amp section of the lock-in, and use toggle switches to configure its "+" input to be ac-coupled, and its "-" input to Ground. Send your desired waveform into the "+" input, and set the pre-amp gain to 1. Now look with your 'scope at the Pre-Amp's output, and you'll see the same waveform as before. But, since this signal is an ac waveform with dc-average value of zero, it's easy to amplify it by a desired factor. Use the rotary switch to get a gain, called here G_1, of 2, or 5, or more. Confirm with your 'scope that the signal is bigger. (Of course, the noise gets amplified right along with the signal!)

2. Now send the Pre-Amp's output to the Filter section's input, and set the Filter to "Q=5". Use the 'scope to look at the Band-Pass output of the filter section, and set the Filter's Frequency Range switch and dial to about the value of the frequency you've chosen for the oscillator that's driving the whole system. Now *tune* the filter to match the (fixed) Oscillator frequency, with correct tuning revealed by getting a maximum strength of signal at the Filter's band-pass output.

The Filter section by itself amplifies the signal by gain G_2, where (for correct tuning) the value of G_2 is given by the Q-value you've selected for the filter. But the filter does *not* amplify all the noise to this same degree – most of the noise is at frequencies *out*side the passband of the filter. So especially when looking at the signal derived from $\Delta V_{out}(t)$, you should be able to see an improvement in the ratio (signal amplitude)/(size of noise) at the filter's output, compared to that ratio prior to the filtering action.

3. Your signal is still an ac signal, occurring at the frequency you picked for the Oscillator. Now send the Filter section's band-pass output to the Lock-In/Amplitude Detector (LIAD) module's Signal input. This module is where the actual process of lock-in detection occurs, so this module is the one that also requires the reference input. Set this module to dc-coupling. To give this module the proper reference input, go back to the Oscillator module's right-hand output, send it to the Phase-Shifter's input, and send the Phase-Shifter's output to the Reference input of the LIAD module. (You can now see why TeachSpin supplies all those short thin black BNC-equipped jumper cables with its lock-in.)

So now the LIAD is getting an amplified and filtered version of the signal at one input, and a reference signal (of the correct frequency, and of a user-variable phase) at the other input. Use a toggle switch to configure the LIAD to its Lock-In mode, and set its gain (here called G_3) to a suitable value – of order 100 for the weak $\Delta V_{out}(t)$ signal. Use your 'scope to look at the output of the LIAD, and vary G_3 until this shows a waveform of height in the range 4-10 V.

The character of this output is that of the input, multiplied by gain G_3, and further multiplied by a factor of ±1 derived from the instantaneous *sign* of the Reference input. As a result, the shape of this waveform depends on the choice of *phase* in the Phase Shifter. Using its Phase and Fine Phase adjustments, get the LIAD's output to look like the lower trace below, taking on only *positive* values:

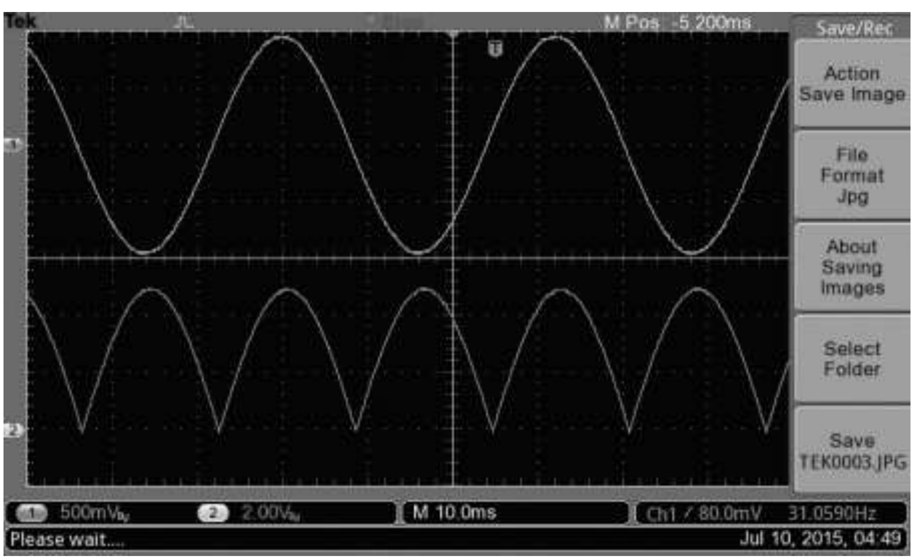

Upper trace: the 'cause' waveform $V_i(t)$; Lower trace: the LIAD module's output, with phase set correctly

Be sure to use *dc*-coupling at your 'scope's input to view this ch. 2 waveform. The entire point of the LIAD module is to accomplish a sort of 'synchronized full-wave rectification' of the signal input, so as to turn an input sinusoidal input, of average value zero, into an output waveform whose minima lie near zero, but which has on the whole a *non*-zero and positive average value.

When you've achieved this, make successive 90° changes in the phase adjustment, and confirm that you can get a LIAD output signal with average value zero, or even negative! Lock-in detection is famously 'phase-sensitive detection', and by mis-adjusting the phase settings of a lock-in you can get the *wrong* sign, or *incorrect* magnitude, of the signal.

4. To be sure you have gotten the phase adjustment just right, temporarily send the strong waveform $V_i(t)$ into the lock-in's input, and scale down the G_1 (pre-amp) and G_3 (LIAD) gain settings until the LIAD output again shows a waveform of height in the range 4-10 V. Now adjust the phases until the result looks like the picture above, and *not* like the results shown below, obtained from slightly erroneous settings:

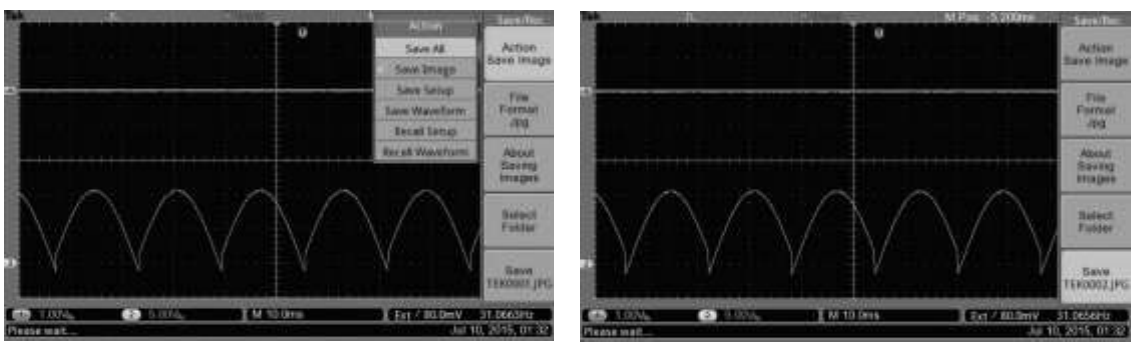

Left image: phase-shift setting in error; Right image: again, a phase-setting error, but of opposite sense

When you have found the right phase setting for the strong and rather noiseless $V_i(t)$ waveform, it will also be the correct phase setting to use for the weaker and noisier $\Delta V_{out}(t)$ waveform, so you *won't need to guess, or to change*, the phase settings hereafter. (You *need* this sort of phase-setting capability, because someday you will use a lock-in in an experiment where the signal is so noise-afflicted that its phase is not even *visible* on a 'scope!)

5. Now send the LIAD's output to the Low-Pass Filter/Amplifier module's input. In this module, set a 'time constant' or averaging time to $\tau = 0.1$ s, and set this module's gain (here called G_4) to value 1. You may set the three-position toggle switch for DC Offset to its central, 0, position; you may also use the '-12 dB/octave' switch setting for best noise rejection. You should now see the ±10-V analog panel meter take on a positive, and approximately steady, value. It's the time-averaged value of the LIAD's output, a running-average taken over of a time-scale of duration τ (and thus over *many*

cycles of individual oscillations of the 'cause' waveform). If the meter reading is near zero, raise the gains G_1 or G_3; if the meter is 'pinned' offscale, lower the gains G_1 or G_3. Once the meter is properly on-scale, its dc reading V_{dc} is a direct measure of the amplitude of the ac signal that you have sent into the Pre-Amp. [Check this, by temporarily setting the toggle switch at the Pre-Amp's "+" input to its "Ground" position instead – the meter should go to zero. (Why?)] The connection between the lock-in's dc output V_{dc} and the input signal's amplitude A is given by

$$V_{dc} = G_1 G_2 G_3 G_4 \cdot C \cdot A \ .$$

Here G_1 = the gain selected in the pre-amp

G_2 = the filter gain, given (at correct tuning) by the Q-value selected

G_3 = the gain selected in the LIAD module

and G_4 = the gain selected in the LPFilter/Amplifier module.

Finally, C is a lock-in constant, whose value ought to be $2/\pi$ (for the detection of purely-sinusoidal signals), but whose value will in fact be unneeded (see below).

To use this lock-in appropriately is to select (among other things) suitable values for the four gain settings. The appropriate values depend on the size of the input signal. For measuring the rather large signal $V_i(t)$, you might try $G_1 = 2$, $G_2 = 2$ (by lowering the Q of the Filter to 2), and $G_3 = 2$ as well. For measuring the weaker $\Delta V_{out}(t)$ signals, you can raise G_1 to 5 or 10, you can revert to $G_2 = 5$ (by returning the Q of the Filter to 5), and raising G_3, perhaps of order 100, until you get a suitable V_{dc} on the meter. We also suggest keeping G_4 as low as possible, perhaps set to 1 or 2. [While commercial lock-ins might include an Overload indicator, here there is nothing to keep you from getting fallacious readings if you choose gains that are too high. So use a 'scope to monitor the preamp and LIAD output waveforms, and see for yourself what happens when the choice of excessive gain causes saturation of the electronic limits of linearity of earlier stages of the amplification chain.]

It is only to be expected that you will need *different* gain values to deal with the different sizes of the waveforms $V_i(t)$ and $\Delta V_{out}(t)$, and you are free to change these settings, provided that with each V_{dc} meter value you record, you also write down the gain settings you used to attain it. But while you might be changing gain settings, you will want to keep *fixed* the Oscillator frequency, the Phase-Shifter adjustments, and the filter tuning, so as to treat equally the 'cause' and 'effect' sinusoids whose amplitudes you are in the process of measuring.

If you aren't content with the ≅2-significant-digit number for V_{dc} that you can read from the panel meter, you may connect, to the LPF/A module's output connector, a DMM serving as dc voltmeter to record the value of V_{dc} more precisely.

Finally then, if for the $V_i(t)$ current-surrogate waveform, you have used settings G_1, G_2, G_3, and G_4, and you have recorded a final dc value of $V_{dc}[i]$, you can infer that the input amplitude was

$$A_i = \frac{V_{dc}[i]}{G_1 G_2 G_3 G_4 \cdot C} \ .$$

And if for the weaker $\Delta V_{out}(t)$ Faraday-rotation waveform, you used values G_1', G_2', G_3', and G_4', and you have recorded a final dc value of $V_{dc}[out]$, you can infer that the input amplitude was

$$A_{out} = \frac{V_{dc}[out]}{G_1' G_2' G_3' G_4' \cdot C} \ .$$

Then the amplitude ratio A_{out}/A_i that you need to get the Verdet constant is given by

$$\frac{A_{out}}{A_i} = \frac{V_{dc}[out]/G_1' G_2' G_3' G_4' \cdot C}{V_{dc}[i]/G_1 G_2 G_3 G_4 \cdot C} \ .$$

Notice that the lock-in constant C cancels out in this ratio – that's the point of using the *same* lock-in to measure $V_i(t)$ and $\Delta V_{out}(t)$ in succession. So then the necessary amplitude ratio is

$$\frac{A_{out}}{A_i} = \frac{V_{dc}[out]}{V_{dc}[i]} \cdot \frac{G_1 G_2 G_3 G_4}{G_1' G_2' G_3' G_4'} \ .$$

Recall that from this ratio of amplitudes, the Verdet constant follows:

$$V = \frac{A_{out}}{A_i} \cdot \frac{1\ \Omega}{V_{max} \cdot k\ L} \ .$$

This is a very sensitive method for finding the Verdet constant V, but it is rather indirectly derived. As a reality check, you might now value all the more those more direct, though less sensitive, initial estimates you made for its value by the methods of Sections 1-3 above.

With these lock-in methods, it's feasible to quantify sinusoidal waveforms of amplitude under 1 μV, and to do so with a remarkable degree of immunity from noise. If you see fluctuations in $V_{dc}[out]$, arising from noise in the weaker Faraday-rotation signal, you can lessen them by choosing a longer averaging time, or time-constant, in the final module. If you do so, remember that you'll need to wait about $5 \cdot \tau$ after any change before a meter reading will stabilize to within 1% of its final value. Confirm this sort of time delay, when using $\tau = 1$ s and observing $V_{dc}[out]$ on the meter, when you have blocked the laser beam near the source with a card, and suppressed any Faraday-rotation signal. Then watch the meter as you unblock the beam, and see that it takes *time* for its asymptotic approach to a new level.

To gain confidence in this lock-in method for measuring a Verdet constant, you can check to see if it gives consistent results when you change the PAA gain, and thus the size of the current in the solenoid. You should be able to explore this variation over about 2 orders of magnitude! Or, at the cost of some re-tuning, you can check to see if the method gives consistent results when you change the oscillator frequency that you use.

To get an appreciation for what lock-in detection has now done for you, recall that your first dc-current detection of Faraday rotation used solenoid currents of order 3 A, and displayed Faraday-rotation angles of a few degrees. What is the smallest current-amplitude at which you have now detected Faraday rotation by lock-in methods? How big (or small!), in degrees, is the Faraday rotation under these circumstances?

5. The Becquerel theory for the Verdet constant

One viewpoint on the Verdet constant of a material is that it's just another 'brute fact' about the substance, akin to its mass density or dielectric constant, a number to be measured empirically and tabulated somewhere. But various observations about Verdet constants do motivate some deeper thinking. For example,

- most materials display a dependence of *V*'s value on the wavelength of light used to measure it, and most materials display *V*(blue) > *V*(red), ie. a Verdet constant larger for higher-frequency light;
- comparing within families of materials, such as the many kinds of optical glass, one finds a distinct correlation between materials of high dispersion (ie. large wavelength dependence in the index of refraction) and high Verdet constant.

These considerations suggest a connection between the Verdet constants and the index of refraction, $n = n(\lambda)$, of a material. Here's how that connection works out.

First of all, there is the famous Fresnel construction, first applied to the earlier discovery of optical activity, which asserts that linearly polarized light can be regarded as a superposition of left- and right-circularly polarized light. If for any reason (such as the chirality or 'handedness' of the molecules in it) a medium exhibits two distinct indices of refraction, n_L and n_R, for these two circular polarizations, then, after propagating a distance L in the medium, light of an initial linear polarization will emerge with a *rotated* linear polarization. The angle of rotation can be shown to be

$$\theta = \pi \frac{L}{\lambda} (n_R - n_L) \, .$$

Note that θ is predicted to scale with L, and notice too that the pre-factor $\pi L/\lambda$ can be enormous (it's of order half a *million* for red light traversing your glass-rod sample), so even *tiny* differences in the two indices, $n_L - n_R$, can cause substantial polarization rotations.

But glass and water are *not* chiral media; yet the presence of a magnetic field seems to make them 'optically active'. Why? Henri Becquerel (of radioactivity fame) was the first to devise a theory for this process. In the era of the newly-discovered electron, he made a classical calculation treating atoms as electron 'oscillators'. We moderns find it easier to translate this idea to a quantum-mechanical model of two-level atoms, and the simplest such models have a spherically-symmetric (s-state) ground state, and have p-states as optically accessible excited states.

For a medium to be transparent in the visible region of the spectrum, the lowest excited states have to lie more than 3 eV above the ground state; a sample of such atoms would exhibit resonant absorption for ultraviolet light (of frequency f, with $h\,f = E_p - E_s$). But absorption and dispersion are necessarily connected; so a medium composed of such atoms has an index of refraction that can be modelled as well, with an index of refraction $n = n(f)$ which exceeds one even in the visible region, and which rises with frequency (toward a singularity at the resonance in the ultraviolet). This agrees with the ordinary dispersion of transparent refractive materials, giving n(blue) > n(red).

Better still, in such a quantum-mechanical model, we can *also* understand the effect of a magnetic field B on the atoms' energy levels. To first order, it leaves the ground s-state unaffected, but it Zeeman-splits the excited p-states into three sublevels separated by two energy splittings, each of size $\Delta E = \mu_B \cdot B$ [where μ_B is the Bohr magneton, $e\hbar/(2m)$]. The two circular-polarization components of light propagating parallel to $\boldsymbol{B}$, labelled as σ_+ and σ_-, will now undergo their resonant absorptions at two new and shifted frequencies:

$$hf_+ = E_p + \mu_B B - E_s \;\; ; \;\; hf_- = E_p - \mu_B B - E_s \; .$$

Not only will the absorption curves differ for σ_+ and σ_- light, so too will the dispersion curves. The same model predicts the effects of the field B on the indices of refraction, and the result is $n_+ - n_- \propto B$, so indeed there should be an optical rotation, proportional to the field strength. In fact, the details of the proportionality can be worked out; there's a fine treatment in the book *The Art of Experimental Physics* by Preston & Dietz. The result of this atomic-physics model, and Fresnel's construction, is Becquerel's prediction for the Verdet constant,

$$V = \frac{-e}{2mc} \cdot \lambda \frac{dn}{d\lambda} \; .$$

Now just as real atoms exhibit the 'anomalous' Zeeman effect, with energy shifts and splitting patterns more complicated than those assumed in the model above for the 'normal Zeeman effect', so too real materials display a more complicated Verdet constant. To a surprising degree, these complications can be captured by a material-dependent (but wavelength-*in*dependent) magneto-optical constant γ, which is dimensionless, and expected to be of order 1. So the revised prediction is

$$V = \gamma \frac{-e}{2mc} \cdot \lambda \frac{dn}{d\lambda} \; .$$

This is a very valuable result, even if the value of γ for a given material has to be fixed by an empirical measurement. That's because spot values of $n(\lambda)$ are readily measured for a given material, and such empirical data on refraction establish an $n(\lambda)$ model (see examples below). From such models, the dimensionless function $\lambda \cdot dn/d\lambda$ follows. Combined with the (universal!) value of $e/(2mc)$ [which, you can show, has value and units 293 rad/(T·m)], the only factor left unspecified is γ. So it follows that measuring V for a material at even <u>one</u> wavelength (as you now have done) allows you to use the model to predict its Verdet constant $V = V(\lambda)$ at *all* wavelengths.

The connection of Verdet constants to index-of-refraction functions motivates a more general discussion of $n(\lambda)$ models. Any optical glass worthy of an identifying number is a carefully-specified and highly-reproducible commercial product, and its refractive-index information is available in tabular form. For example, SF-57 glass has an index of refraction carefully measured (to 6 significant figures!) at the wavelengths of certain spectral lines. Published results are

wavelength λ (in μm)	index *n*
0.4358	1.89391
0.4800	1.87425
0.4861	1.87205
0.5461	1.85504
0.5893	1.84636
0.6328	1.83957
0.6438	1.83808
0.6563	1.83651
0.7065	1.83104
0.8521	1.82045
1.0600	1.81184

Plotting such data reveals a smooth dependence of index *n* on wavelength λ. A very early model for such *n*(λ)-dependence was Cauchy's model,

$$n(\lambda) = A + \frac{B}{\lambda^2}\ .$$

You can plot the data tabulated above in a graph of *n* vs. $1/\lambda^2$, and look for a nearly-linear dependence; from such a plot you can extract Cauchy *A*- and *B*-parameters for SF-57. But given such an *n*(λ) model, it's easy to compute the function $\lambda \cdot dn/d\lambda$ ($= -2\,B/\lambda^2$) and make a decent prediction of its value at any given wavelength (such as 0.65 μm).

Your Cauchy plot will reveal a fit *in*adequate to the precision of the data, and to fit properly such 6-sig-fig data at 10 or more wavelengths takes more parameters, and a more physically-motivated fitting function. One model often used in optics is called a Sellmeier expansion:

$$n^2(\lambda) = 1 + \frac{B_1\lambda^2}{\lambda^2 - C_1} + \frac{B_2\lambda^2}{\lambda^2 - C_2} + \frac{B_3\lambda^2}{\lambda^2 - C_3}\ .$$

where $n^2 - 1$ is being fit by a multiple-resonance model. (Clearly the model would exhibit singular behavior at wavelengths $\lambda = \sqrt{C_1}, \sqrt{C_2}, \sqrt{C_3}$.) A typical model uses two such UV-resonances and one IR-resonance. Six coefficients that capture *n*(λ) [for λ expressed in μm] for SF-57 glass are

B_1	1.8165 1371
B_2	0.4288 9364
B_3	1.0718 6278
C_1	0.01437 04198
C_2	0.05928 01172
C_3	121.419 942

Here too you can compute a model function for $n(\lambda)$, you can test it against the tabulated $n(\lambda)$-data for adequacy, and (more to the present point) you can use it to compute $\lambda \cdot dn/d\lambda$ at any desired wavelength.

That model of $\lambda \cdot dn/d\lambda$, combined with your 'spot measurement' for $V(\lambda = 0.65\ \mu\text{m})$ for SF-57 glass, will give you a value for the constant γ. And with that constant in hand, you can predict (without any further 'fudge factors') how the Verdet constant V ought to depend on λ over the entire visible spectrum. Make yourself such a predictive plot! And if you have access to other wavelengths of mW-level laser beams, you can test your own prediction.

6. The 'effective length' in the case of long samples

The use of lock-in detection makes your Faraday-rotation apparatus sufficiently sensitive to enable the detection of Faraday rotation in samples much more ordinary than exotic high-index glass such as SF-57. You could try samples made of any available glass or even plastic, or you could use samples of shorter length, and still detect the Faraday effect.

But to make even more samples accessible to you, TeachSpin supplies a glass cell, to be filled with any liquid (transparent at 650 nm), to permit the measurement of its Verdet constant. The optical and electronic techniques you'd use are identical; the only complication is that the cell is *too long* to be treated as if it placed all of the sample in a constant-field region near the center of the solenoid. So this section teaches you how to cope with the spatial variation of the magnetic field B along the length of an axis along the light beam. Formerly we took B as constant, and having value $k \cdot i$, for the full length L of the (glass) sample. Now we'll use coordinate z along the axis of the solenoid, and assume the solenoid gives a field variation with z, and we'll write that 'field profile' by assuming

$$B \rightarrow B(z) \equiv B_{center} \cdot \frac{B(z)}{B_{center}} = k\, i \cdot f(z)\,.$$

So by definition $f(z)$ has value 1 at the center of the solenoid, and values just below 1 in its central section; but the f-values drop to about ½ at the ends of the solenoid's windings, and drop farther still beyond those ends. Note that $f(z)$'s shape could be computed analytically from the theory of solenoids, or it could be obtained from 'surveying a map' of $B(z)$-values, obtained (for example) with a Hall-effect probe. But what would we *do* with $f(z)$ profile information, to enable the computation of a Verdet constant for a long sample? Here's the use of it:

Formerly we assumed Faraday rotation $\theta = V\,B\,L$ for a sample of length L, and clearly that result specializes, for length dz of sample, to rotation $d\theta = V\,B\,dz = V\,B(z)\,dz$. But using our field profile information, we can write that as $d\theta = V \cdot k\, i\, f(z) \cdot dz$. Now we integrate over the entire region of the sample $\{S\}$ to get net Faraday rotation

$$\theta = \int_S d\theta = \int_S V \cdot k\, i\, f(z)\, dz = V \cdot k\, i \int_S f(z)\, dz\ .$$

Formerly, we assumed $f(z)$ was constant and of value 1 over the length L (=0.100 m) of the glass sample, so this integration gave

$$\theta = V \cdot k\, i \int_S 1\, dz = V \cdot k\, i \cdot L\ .$$

Now instead we get

$$\theta = V \cdot k\, i \int_S f(z)\, dz \equiv V \cdot k\, i \cdot L_{eff}\ .$$

so clearly all of our old results for data reduction hold true, provided we replace references to L by L_{eff}, where L_{eff} is the 'effective length' of the sample, given by

$$L_{eff} \equiv \int_S f(z)\, dz \ .$$

Notice the limits of integration are over the length of the sample S (*not* over the length of the solenoid). Notice that $f(z)$ is dimensionless, with value 1 at the solenoid's center, so L_{eff} has dimensions of length arising from the integration of dz.

> [If a sample were to be *really* long, extending far beyond both ends of the solenoid, the integral we need is related to one from Ampere's Law, since for an N-turn solenoid we'd have
>
> $$\int_{-\infty}^{\infty} B(z)\, dz = \oint_C \boldsymbol{B} \cdot \boldsymbol{dl} = \mu_0\, N\, i \ .$$
>
> The Amperean path C has to be 'closed at infinity' for this to work, but a 'semicircle at infinity' gives zero contribution, and the whole of a closed-path integral is thus given by an integration on an infinite line, our z-axis. In our model for field profile, this gives
>
> $$\int_{-\infty}^{\infty} k\, i\, f(z)\, dz = \mu_0\, N\, i\,,$$
>
> from which we could extract
>
> $$L_{eff} = \int_S f(z)\, dz = \frac{\mu_0\, N\, i}{k\, i} = \frac{\mu_0\, N}{k}$$
>
> for a really long sample. In the case of the TeachSpin solenoid, this would give
>
> $$L_{eff} = \frac{\mu_0\, N}{k} = \frac{(4\pi \times 10^{-7}\,\text{T·m/A})(1400)}{0.0115\ \text{T/A}} = 0.1530\ \text{m} \ .$$
>
> In practice, the liquid-sample tube does not extend so far as $\pm\infty$, and thus its value for L_{eff} will be shorter than this, and certainly shorter than the tube's own physical length. (Notice that the number for L_{eff} in the very-long-sample limit differs from the physical length of the solenoid, 6″ or 0.1524 m. Note also that it's the $k{\cdot}L_{eff}$ product that you would need to extract a Verdet constant from your data, and in the very-long-sample limit, that product reduces to $\mu_0\, N$, which requires no metrology except turn-counting!)]

You should first try this experiment using ordinary (or distilled or de-ionized) water as a sample. You'll need less than 15 cm^3 of liquid to fill the tube. A syringe with the appropriate hypodermic needle is a good tool for filling the tube. [To empty the tube, the injection of air by syringe into the tube to displace the liquid is fairly effective, though this will leave behind some liquid on the interior of the tube. A continuing flow of oil-free dry air, delivered into the inside of the sample cell via a thin tube, will eventually evaporate the last of the water.]

You can also try other liquids, *provided* you can plan ahead on how to remove them safely from the inside of the tube. You can also try water solutions of various chemicals; solutions of paramagnetic salts are particularly interesting, since they permit you to detect the change in *sign* of the Verdet constant with increasing concentration of salt.

Finally, there's the issue of a 'control group'. You might suppose that for an experiment on the Verdet constant for water, you should compare

(Faraday rotation with a water-filled tube present)

with (Faraday rotation with a water-filled tube absent).

But the *properly* controlled experiment is to compare

(Faraday rotation with a water-filled tube present)

with (Faraday rotation with an empty tube still present).

There is no guarantee that this last, empty-tube, investigation should give zero Faraday rotation! That's because an empty tube does put into the light beam two samples made of glass (heretofore thought of as 'end windows'), and it puts those samples into regions of weak, but non-zero, magnetic fields beyond the ends of the solenoid. See if the sensitivity of your lock-in method allows you to detect this 'fringe field' effect, or to put an upper bound on it. The lock-in method is certainly sensitive enough for you to detect the Faraday rotation caused by a water-filled tube, so you will be able to measure the Verdet constant for water.