

The Anarchist Library
Anti-Copyright



Manipulation As Withholding

William Gillis

19th February 2016

One of the most disappointing things about the anarchist community is that while we've widely recognized that our critique of power relations extends to interpersonal dynamics we've — so far — largely shied away from addressing such in concrete terms. One-on-one we often repeat accusations and condemnations of manipulation, but we've never sorted out precisely what manipulation is. Even in the tumblr renaissance there's a marked tendency for the heroic young folk decrying power structures within various scenes or milieus to suddenly shy away when it comes to identifying and arguing against manipulation *itself*. Instead we generally mutter about how such and such example of manipulation rhymes with aspects of certain widely recognized systems of oppression. But interpersonal power dynamics are not bad merely because they *reflect* currently widespread shitty systems. Bespoke patterns of abuse are still abusive, regardless of whether or not we can easily shove them into categories like “racist” or “patriarchal”.

In one of the blessed signs that we *are* trying to work some of this out there's been an increasing amount of explicit and sharp disagreement over what constitutes “gaslighting.” When one side argues that gaslighting is not treating someone's beliefs as innately

William Gillis
Manipulation As Withholding
19th February 2016

[http://humaniterations.net/2016/02/19/
manipulation-as-withholding/](http://humaniterations.net/2016/02/19/manipulation-as-withholding/)

theanarchistlibrary.org

valid and the other side argues that it's making/letting someone become disconnected from objective reality it's quite clear that there's a deep divide in philosophical foundations going on. The sort of hugely substantive split that we absolutely cannot afford to paper over but desperately need to drag into the light and hash out.

I've applauded these debates, but I want to go further and argue that there's a grave philosophical chasm to be found in our definition of "manipulation" itself.

I think this Metafilter exchange teases the fracture apart quite well:

Person 1:

Please do not so as I request if there is any taint of fear or punishment if you don't.

Please do not do as I request to buy my love, that is, hoping I will love you more if you do.

Please do not do as I request if you will feel guilty or shameful if you don't.

And certainly do not do as I request out of any sense of duty or obligation.

Person 2:

I don't think I'd ever do much of anything without these motivating factors.

I want to be absolutely clear: I am not advocating a nihilism on the subject of manipulation; I think that there exists a clear and concrete definition of manipulation that is quite aggressively expansive and deserves unqualified condemnation and resistance from anarchists. Person 2's response can be read in a very nihilistic manner about motivation, and I'm certain many sociopaths & social capitalists would instinctively read it as agreeing with them that *"literally everything is manipulation, it's just that some people*

consider some random types of it objectionable.” I am not taking that position, but I do think the response correctly highlights just how extreme the first person’s framework is.

What’s implicit in the above list of proclamations is a condemnation of anything that’s not a direct and immediate personal desire. All the motivations presented as invalid or to be avoided — fear of negative reactions, hope to deepen love, guilt, obligation — are united by involving higher-order desires or considerations. Sure you may happen to sometimes feel an immediate desire, but these desires are often overruled or outweighed, upon rumination, by other more deeply rooted desires.

In a direct reading of this philosophy if I make a rational and honest argument appealing to your conscience — telling you that if you don’t stop playing video games and call an ambulance for me you’ll feel guilty for going against a more underlying desire when I bleed out — that would be classified as “manipulation.” But surely if you actually do deeply care about whether I live or die it’s hardly manipulative for me to help you remember that while you happen to lie in a temporary video game coma. The other terms have equally absurd implications under this frame of mind, and the reductions for them should be obvious. Are we never to act in ways we hope will deepen someone’s love for us?

You might find this a particularly uncharitable read, but there’s a long tradition of thought in the Left that breaks in this direction. An example would be the recurring claim by many Foucauldians that persuasion of any kind — including rational appeals to one’s underlying desires or values and the ramifications from them — is coercion. This syncs up well with a current of immediatism found within the ‘radical’ milieu and many youth subcultures that instinctively takes a very hands-off approach to the roots and ramifications of one’s feelings and calls that ‘autonomy.’ This is the “*Thinking About One’s Desires Is Bad*” philosophy that idealizes desires (like love, etc) as mere magical weather patterns that just happen to strike us whenever, and that we dare not examine closer or seek

to clarify lest our understanding and agency in them diminish their ‘wildness.’ It’s closely tied to the notion that thinking about shit is dangerous because it leads to sociopathy and what differentiates us from sociopaths is *not knowing*, treating other people’s minds (and our own) as black boxes.

There is a very profound question here that professional philosophy as a whole has unfortunately largely ducked: which is what if you *can* actually map and predict someone’s responses with great accuracy, what are your ethical obligations? If you come to understand someone to the point where they cease being a black box? If interactions with them are functionally deterministic on a certain scale or within a certain context (however limited)? At that point whatever action (or inaction) you take, you’re determining what ends they move towards. This is the sort shit that creates nihilists/fascists — and non-anarchists more generally — because they instinctively define power relations / manipulation in terms of causal interplay, and since you can’t get outside causality you can’t get outside power / manipulation.

A prototypical liberal response is to react by clinging to ignorance. To say “*well you don’t know!*” but also “*you shouldn’t know!*” To basically forbid exploration. In this light “manipulation” is taken to be any situation where one person doesn’t treat the other’s thoughts or feelings as a black box. In the extreme limit this ends up rejecting things like, “hey, I know that you like baseball cards so you should probably go to the baseball card festival they’re putting on.” Or, “*I feel safer with partners who are honest [and thus being honest with me will likely heighten my capacity to love you].*” Both actions leverage knowledge of a person’s internal desires and contextual ramifications to change their likely behavior, even their desires. If this is manipulation then rejecting it means rejecting knowing people and/or any deduction of ramifications from desires, much less prioritization of desires.

default, you can freely persuade, you can freely argue, you can freely influence other people. But you have to be proactively honest. Telling them not just the stuff that cuts in your direction, but — where possible — the stuff that goes against, and generally all the stuff that might be remotely relevant to them. You have to push back against any common cognitive fallacies they might be otherwise subject to — including the influence of often unconscious dynamics like your own social standing or creeping halo effect.

The world is a complicated messy place, we never know anything absolutely, much less precisely what others would be better informed and empowered with greater agency by knowing, but the complexity of the world is no excuse for failing to engage with it. And in many situations we *can* judge or know these things quite well.

Of course actual manipulation in the real world typically involves a vast contextual network built from the sort of primitive dynamics described here, and often interweaves in subtle ways with contexts of physical coercion, and we might want to classify hardwired processes like some forms of PTSD that one has no conscious agency in as directly *physical* forms of control, but the point remains that proactive honesty would make manipulation itself impossible, even if other forms of control remain to be tackled.

Now I'm not suggesting that our world isn't a complicated place — or less euphemistically a horrific dystopia — where utilizing power relations isn't sometimes called for, even by those of us committed to their ultimate abolition. Punching a cop in the face is, obviously, a power relation. And privacy, secrecy, etc are quite often functionally necessary, especially in movements of resistance. There can be good reason to avoid letting our boss know about our politics or even our friends know about all our kinks, if only for politeness' sake. Similarly while the hierarchies of information access that drive ingroup/outgroup dynamics are corrosive to rationality or intellectual vigilance, they can be quite useful for the oppressed or in defensive situations more broadly. There are thousands of diverse ways in which navigating even basic human interactions involve manipulation. But anarchists are not marxists, enthusiastic or untroubled by embracing power as a means, we believe in ferreting out even the most subtle of power dynamics, and that while ends and means are deeply interconnected and we cannot and should not act entirely as saints, we should nevertheless strive to build as best we can an image of the ends we would like to see. Namely, a world without power relations.

While strategy and pragmatism are surely called for, our lives must be coherent with that ultimate goal or value. And that means a pressure or general tendency against manipulation in any form, barring limited and painfully necessary exceptions. But yes, as a

I want to be fair: such an expansive notion of “manipulation” wouldn't necessarily forbid literally *all* human interactions, but it would obviously dramatically constrain them. Far more than I think our natural concept of and objection to “manipulation” can defend.

This is usually where the average person stops in their private analysis. And it provides a good representation of the ignorance-as-egalitarianism vs intelligence-as-sociopathy assumed dichotomy that I've long been emphasizing underpins so damn much of people's philosophies and our world's institutions and ideologies. But no, I think there's a perfectly functional way to define manipulation that doesn't go to such absurd lengths, rejection of which is perfectly compatible — even convergent — with at the same time valuing intelligence and understanding.

Manipulation is *withholding* information from people in order to constrain — and thus presumably channel — their choices.

For example if you have an argument that (alone) would persuade someone to do B, but you *also* have an argument that might persuade them to do A, and yet you only provide them with one of the arguments then you have manipulated them. Similarly even if you haven't *found* another argument beyond the one you're presenting, if you're working off an array of remotely relevant context that they're not privy to and you fail to present the full context you have (so they might search within it for new or better arguments), you're manipulating them.

Under this definition a condemnation of manipulation is straightforward in its prescription: if you can model someone in a given context quite well the way to respect their agency is to elevate them to whatever level you're on (at least in all relevant respects).

Let's quickly state this more rigorously for the nerds:

Take for instance two AIs that are exactly equal, if one reaches a computational conclusion earlier than the other it should be free to transmit that proof to the other, even though such a transmission

utterly determines the result / course of action the other AI will take. However if one AI is larger than the other AI and can thus model it internally, the extra bits of that AI functionally represent an oracle. The ethical obligation on the AI is to augment the lesser AI with said oracle. It doesn't matter what these extra bits are — an extra level of recursion, extra information about the environment, etc. The course of action by the bigger AI that *doesn't* count as manipulation of power dynamics is to help the smaller AI enhance itself. In short: to expand agency.

Of course living breathing human beings are incredibly complicated and messy, with all kinds of subjective barriers to personal experience. No one can ever perfectly simulate another human within their own head. However there clearly are situations where this is possible *to some degree* — and the course of action that *isn't* manipulative is the one that shares whatever information that might be relevant to another person's life and thoughts and decisions. Which applies just as strongly for two people of more or less equal intelligence but unequal sets of knowledge.

This has been my guiding philosophy for over a decade.

Some context might be illustrative as an example (and perhaps fend off accusations of disconnect or resentment):

I tried to navigate high school with a moral code, refusing to do a huge array of things, and yet I accumulated some fair amount of social power and standing nonetheless. My prohibitions were far more restrictive than anyone else's I've ever met. And yet I still found myself functionally manipulating people to vast degrees. One easy to explain mechanism is that I was a lot of people's counselor and confidant — in part because I was compassionate and helpful but also because I kept their secrets. Now I never spilled people's secrets to hurt them or discredit them or make power plays — and I wouldn't have. But I nevertheless *had* said secrets. This means two things, 1) a number of people still had some lurking concern that I might spill them if they pulled shit with me, but

more importantly 2) I had in many respects a better map of the world than everyone else. I knew the social landscape.

If I know that person A's crush on person B is reciprocated, I suddenly know a lot more about the possible ramifications of normal casual actions that I might take. I may be sworn to secrecy and non-meddling by both, but when picking who to invite along with a group skipping class I can't pretend that I'm not either increasing their likelihood of hooking up or decreasing it. I am functionally in control whether or not I want to be. The only thing that reduces my control over the situation is to increase the operating knowledge both of them have.

One quickly realizes that more obtuse manipulation through for example lying is but a subset of withholding. Lying that is immediately revealed as lying is merely joking. It is the one-two punch of changing the landscape by introducing new data ("Jane says she wants you to go") but not the relevant completion of that data (the single extra bit that clarifies "this is a lie") that effects manipulation. And joking that depends on a persistent lack of clarity over whether or not someone is joking can be easily seen to replicate power through differentials of awareness. For example teasing that's harmless or convivial between bros who are in on the unseriousness can become functionally abusive when directed at someone who isn't on the "in". Even if the bro doesn't consciously mean to leverage power or ignorantly assumes the person on the out will get it, the information differential creates a power dynamic anyway. A much more evolved form of this is visible in meme ecosystems online today, where the function of in-jokes has been stripped bare and weaponized. New memes — in the sense of in-jokes — are constantly created and dispersed to reveal the lines of association and reinforce hierarchies of information access (that overlap in complex meshes and constantly shift). Similarly, exploiting someone's subrational instincts or cognitive fallacies is almost always a matter of withholding your awareness of such processes and what you consider the likely impact of your actions.