Bayesian Expectancy Invalidates Randomized Controlled Medical Trials*

Gilles Chemla
Imperial College Business School, DRM/CNRS, and CEPR.

Christopher A. Hennessy
London Business School, CEPR, and ECGI.

May 2016

Abstract

The randomized controlled trial (RCT) is viewed as the gold standard in eliminating placebo effects and identifying pure non-placebo physiological effects. Expectancy theory, the leading theory of placebo effects, posits that experimental subjects manifest improved health quality in response to expectations of higher future health quality. We show that if test subjects engage in rational Bayesian updating based upon the direct physiological effect they experience during an RCT, the expectation of placebo effects are, under general conditions, unequal across treatment and control groups. It follows that difference in average outcomes between treatment and control groups is a biased estimate of the direct physiological effect. Biases vary with funding and regulatory constraints, as well as the probability of next-generation drug improvements.

*Corresponding author (Hennessy): Regent's Park, London, NW1 4SA, U.K.; chennessy@london.edu; 44(0)2070008285.

1 Introduction

A critical objective for medical research is to measure the non-placebo physiological effect of a treatment, also known as *characteristic effect* (Grünbaum (1986)) or *specific effect* (Malani (2006)). The randomized controlled trial (RCT) is viewed as the gold standard in measuring non-placebo physiological effects. In fact, describing the rise of randomized trials in medicine in *The Lancet*, Kaptchuk (1998) writes, "The greater the placebo's power, the more necessity there was for the masked RCT itself." This gold standard status of RCTs is codified by the International Conference of Harmonization (2000), ratified by medical regulators in the E.U., U.S., and Japan. In fact, the perceived reliability of the RCT has caused the methodology to be emulated in the social sciences. For example, the RCT is held up by Angrist and Pischke (2009), as the ideal for achieving unbiased estimates of causal effects in the social sciences.

The argument for the randomized, controlled, double-blind medical RCT is by now so familiar that it typically escapes discussion. Health quality is viewed as the sum of the non-placebo physiological effect plus any brain-modulated physiological ("mental") effect. If subjects are randomly assigned across treatment and control groups, and blind to their assignment, then the expectation of mental effects is posited to be equal across treatment and control groups. It would follow then that the difference in average health outcomes across groups yields an unbiased estimate of the expectation of the non-placebo physiological effect.

The traditional proof of RCT validity assumes mental effects are additive identically distributed random variables independent of the assigned group. However, as shown below, the independence assumption is at odds with a rational Bayesian formulation of *expectancy theory*, the leading theory of placebo effects. Stewart-Williams and Podd (2004) state, "On the expectancy account, the effects of such factors come through their influence on the placebo recipient's expectancies." Malani (2006) writes, "placebo effects cause more optimistic patients to respond better to treatment than less optimistic patients." In this spirit, in perhaps the first known placebo-controlled trial, Haygarth (1801) wrote that his study, "clearly prove[d] what wonderful effects the passions of hope and faith,

excited by mere imagination, can produce on disease."

While beneficial brain-modulated physiological effects are traditionally viewed (modeled) as irrational noise terms, we here consider that such effects arise rationally from the expectation of better future health outcomes. For example, expectation of less future pain may reduce stress, improving outcomes today for subjects suffering from ulcers or hypercholesterolemia. Similarly, expectation of higher future survival rates may alleviate the severe anxiety associated with life-threatening diseases, with relaxation, rest and sleep improving health outcomes today. As a final example, expectation of improved future brain functioning, that is, alleviation of hopelessness, may mitigate the physical effects of depression today.

We show that the expectation of mental effects should not be presumed to be equal across treatment and control groups in RCTs. Intuitively, with rational Bayesian test subjects, beliefs will vary systematically with the probability distribution of the respective direct physiological effects. Unless the probability distribution governing direct physiological effects is equal across treatment and control groups, beliefs will differ across these groups, implying unequal mental effects in expectation. The difference in average health outcomes across treatment and control groups then delivers a biased estimate of the direct physiological effect.

There have been few attempts by economists to formally model placebo effects. An important exception is Malani (2006). However, his objective was to develop a formal empirical test for placebo effects. The key insight derived from his model is that if there are indeed placebo effects, then treatment and control groups should manifest better outcomes in trials with higher treatment probability. Strikingly, after confirming this prediction for the treatment arm, Malani states, "An important *footnote* (our italics) to the findings reported in this paper is that regression analysis also reveals a positive correlation between the probability of treatment and outcomes in the control group, but only when the control therapy is active treatment, namely H_2 blockers in PPI trials. It does not find a positive correlation when the control is an inert pill." That is, Malani's empirical findings suggest that placebo effects differ across treatment and control groups in a manner dependent upon

the physiological effect of the control. As shown below, such cross-arm placebo differentials are inconsistent with standard additive placebo effects, such as those modeled by Malani (2006). In this paper we offer a rational Bayesian theory for such differentials, one that calls into question the logical basis for the presumption of lack of bias in RCTs.

Of course there is a large medical literature on RCTs. Rothwell (2005, 2006) provides excellent surveys. Existing critical examinations of medical RCTs have emphasized the difficulties in their practical implementation. For example, Rothwell (2005, 2006) discusses numerous practical sources of selection bias. Further, he discusses how trial protocol will tend to differ from real-world implementation, and also discusses the difficulty in measuring comprehensive health outcomes accurately. Such criticisms apply to the difficulty in implementing RCTs but they do not directly challenge the logical basis of the RCT methodology itself. Consistent with our model, there is also an empirical medical literature documenting discrepancies between RCT outcomes and long-term clinical outcomes, e.g. rheumatoid arthritis osteoporosis, post-myocardial infraction, multiple sclerosis and heart failure (Rothwell (2005), Pincus (1998)). This may not be due only to imperfect implementation of double-blind RCTs, but rather may point to a flaw in the very concept of double-blind RCTs.

2 The Model

The technology is common knowledge. There are two dates $d \in \{1, 2\}$. At $d = 1$, a double-blind randomized, parallel group, controlled trial (RCT) is conducted. The objective of the RCT is to measure the efficacy of a novel drug. The measured efficacy determines the probability of the tested drug being distributed at $d = 2$, as well as the probability of next-generation improvements to the drug. Depending on the setting, one can think of the control group as being given either an inert pill (placebo-controlled trial) or some traditionally-used pill (controlled trial).

In the model, the RCT is ideal, with the test panel being equivalent to the afflicted population,

eliminating self-selection concerns.[1] The afflicted population $\mathcal{I}$ is a measure one continuum of ex ante identical agents, eliminating the concern expressed by Deaton (2010) regarding small sample bias. Below, $\mathcal{T}$ ($\mathcal{C}$) denotes the set of agents randomly assigned to the treatment (control) group. The measure of the treatment group is $t \in (0,1)$.

During the RCT ($d = 1$), agent $i \in \mathcal{I}$ experiences a direct physiological effect p_i^1, a random variable with support $\mathcal{P} \equiv \left[\underline{p}, \overline{p}\right]$. This is the health quality that would be experienced absent any placebo effect. If $i \in \mathcal{C}$, the direct physiological effect is an independent draw from an atomless twice continuously differentiable cumulative distribution F_C, with probability density f_C. If $i \in \mathcal{T}$, the direct physiological effect is instead an independent draw from an atomless twice continuously differentiable cumulative distribution F_S, with probability density f_S. The treatment group distribution function is state-contingent, with the efficacy state $S \in \{L, H\}$.

Let $\lambda \equiv \Pr[S = H]$, and assume $\lambda \in (0,1)$. The expectation of the direct physiological effect is

$$\mu_J \equiv \int_{\mathcal{P}} p f_J(p) dp,\ J \in \{C, L, H\}.$$

The following conditions are imposed.

Assumption 1: *The expectation of direct physiological responses is such that* $\mu_H > \max\{\mu_C, \mu_L\}$. *For each* $p \in \mathcal{P}$ *there exists a* $J \in \{C, L, H\}$, *such that* $f_J(p) > 0$. *For each* $p \in (\underline{p}, \overline{p})$, $f_J(p) > 0$ *for all* $J \in \{C, L, H\}$.

The first part of Assumption 1 ensures there is at least one state (H) where agents would consume the new drug voluntarily at $d = 2$ if it is offered. The second part of the assumption ensures beliefs are well-defined on $\mathcal{P}$. The third part of the assumption ensures the derivative of beliefs is well-defined on $(\underline{p}, \overline{p})$.

During the RCT agent $i \in \mathcal{I}$ experiences health quality Q_i^1, where

$$Q_i^1 \equiv p_i^1 + M_i. \tag{1}$$

[1] See Malani (2006) for a detailed analysis of self-selection in RCTs.

In the preceding equation, the first term, the direct physiological effect, is assumed to be observable to the agent. However, this assumption is not necessary when beliefs are monotone. Under monotonicity, the direct physiological effect can be inferred from the total effect. The second term captures the brain-modulated physiological effect.

Since additivity of the mental effect plays an important role in the traditional proof for the unbiasedness of RCTs, we emphasize its adoption here.

Assumption 2: *The brain-modulated physiological effect enters health quality additively.*

When the state is determined to be S, then with probability π_S the newly-tested drug is again distributed at $d = 2$ and further research into next-generation improvements to this drug are implemented. The benefit of research into next-generation improvements to the new drug is to increase health quality by the increment $\delta_S \geq 0$. We assume that continuation is more likely in state H, with $\pi_H \geq \pi_L \geq 0$. If $\mu_L + \delta_L \leq \mu_C$ we impose $\pi_L = 0$. Assuming, as we do, that agents are risk-neutral, this assumption ensures all agents will take the new drug voluntarily if it is offered at $d = 2$.

For simplicity, assume the effect of both the treatment and control drugs is non-cumulative. In particular, conditional upon S, for each agent $i \in \mathcal{I}$, and for each drug type, the direct physiological effect at $d = 2$ is independent of the direct physiological effect at $d = 1$. That is, p_i^2 is an independent draw of the direct physiological effect from the relevant distribution, specifically F_S if the new drug is offered and F_C if not. Let Φ_S be an indicator for continuation (distribution plus research into next-generation improvements). We have

$$\pi_S \equiv \Pr[\Phi_S = 1].$$

Health quality at $d = 2$ can then be expressed as:

$$Q_i^2 \equiv p_i^2 + \Phi_S \delta_S. \tag{2}$$

Let β denote the probability assessment that $S = H$ based upon the direct physiological effect experienced by the agent during the RCT. From Bayes' rule we have:

$$\beta(p) \equiv \Pr[S = H | p_i^1 = p] = \frac{\lambda \left[t f_H(p) + (1-t) f_C(p) \right]}{t \left[\lambda f_H(p) + (1-\lambda) f_L(p) \right] + (1-t) f_C(p)}. \tag{3}$$

For brevity, let:

$$X(p) \equiv \mathbb{E}[Q_i^2 | p_i^1 = p].$$

We have:

$$X(p) = \beta(p)[\pi_H(\mu_H + \delta_H) + (1 - \pi_H)\mu_C] + [1 - \beta(p)][\pi_L(\mu_L + \delta_L) + (1 - \pi_L)\mu_C]. \tag{4}$$

Two points are worth noting in the preceding equation. First, expected future health quality varies with the continuation probabilities (π_L, π_H). Second, expected future health capitalizes not only the effects of the new drug but also anticipated improvements to the drug, as captured by the improvement parameters (δ_L, δ_H).

Consistent with expectancy theory, we adopt the following assumption.

Assumption 3: *Brain-modulated physiological effects are equal to* $\Psi(X)$, *with* Ψ *being continuously differentiable and strictly increasing on its support.*

Under Assumption 3, the mental component of health quality at $d = 1$ can be computed as:

$$M(p) = \Psi\left[X(p)\right] \Rightarrow M_i \equiv M(p_i^1). \tag{5}$$

Under the stated assumptions, the expected treatment-control health quality difference is:

$$\mathbb{E}[Q_i^1 | i \in \mathcal{T}, S] - \mathbb{E}[Q_i^1 | i \in \mathcal{C}, S] = \underbrace{\mu_S - \mu_C}_{\text{Non-Placebo Effect}} + \underbrace{\{\mathbb{E}[M_i | i \in \mathcal{T}, S] - \mathbb{E}[M_i | i \in \mathcal{C}, S]\}}_{\text{Bias}}, \tag{6}$$

where

$$\mathbb{E}[M_i | i \in \mathcal{T}, S] - \mathbb{E}[M_i | i \in \mathcal{C}, S] = \int_{\mathcal{P}} M(p)[f_S(p) - f_C(p)]dp. \tag{7}$$

It is traditional to think of mental effects, the M_i, as being i.i.d. random variables. Under the traditional interpretation, the bracketed term in equation (6) is equal to zero, implying the average treatment-control health quality difference yields an unbiased estimate of the direct physiological effect of the drug relative to the control. Thus, the absence of bias in RCTs can be understood as being predicated upon two assumptions: additivity and i.i.d. mental effects.

Although not our focus, we note the additivity assumption may be questionable in some contexts. For example, in settings with upper (complete recovery) and lower (death) bounds on the outcome

variable, max and min functions would be applied to $p + M$, so the expectation of mental effects would not necessarily cancel across treatment and control groups even if the M_i were, in fact, i.i.d. random variables.

In order to highlight the main causal mechanism, it is instructive to first derive a condition under which absence of bias is assured.

Proposition 1 *A sufficient condition for the expected treatment-control health quality difference to equal the expectation of the direct physiological effect, regardless of the true efficacy state S, is*

$$\pi_H(\mu_H + \delta_H) + (1 - \pi_H)\mu_C = \pi_L(\mu_L + \delta_L) + (1 - \pi_L)\mu_C.$$

Proof. $M' = \Psi' X'$. Under the stated condition $X' = 0$.■

The intuition for the preceding proposition is as follows. Under the stated condition, expected future health quality is equal across states H and L, implying variation in beliefs (β) across the treatment and control groups is inconsequential, since differences in beliefs then have no effect on brain-modulated physiological effects. Conversely, if expected future health quality does vary across the states, then differences in beliefs across the treatment and control groups are indeed consequential. And differences in beliefs across treatment and control groups arise directly from differences between the distributions governing their respective draws of direct physiological effects.

In light of the preceding proposition, the remainder of the analysis imposes the following additional assumption.

Assumption 4: *The expectation of future health quality is strictly higher in state H than in state L.*

3 Bias in RCTs

We now describe a number of settings in which double-blind medical RCTs are biased. In each setting, the belief function β is strictly increasing, implying the mental effect function M is also strictly increasing. That is, the technology is such that a better draw of p during the RCT causes the

agent to assign a higher probability to state H, which leads to better brain-modulated physiological responses. With M demonstrated to be increasing, we use stochastic dominance relationships across the distribution functions to determine the sign of the bias.

For brevity, let

$$R_S(p) \equiv \frac{f_S(p)}{f_C(p)} \ \forall \ S \in \{L, H\} \text{ and } p \in (\underline{p}, \overline{p}).$$

Differentiating β and rearranging terms one finds:

$$\text{Sign}[\beta'(p)] = \text{Sign}\left[\frac{[tR_H(p) + (1-t)]'}{tR_H(p) + (1-t)} - \frac{[tR_L(p) + (1-t)]'}{tR_L(p) + (1-t)}\right]. \tag{8}$$

Using the preceding equation, we have the following proposition.

Proposition 2 *If $\frac{f_H}{f_C}$ and $\frac{f_C}{f_L}$ are both strictly increasing (MLRP), then the expected treatment-control health quality difference is greater (less) than the expectation of the direct physiological effect in state H (L).*

Proof. Under the stated conditions it follows from equation (8) that β is strictly increasing on $(\underline{p}, \overline{p})$, from which it follows M is also strictly increasing on $(\underline{p}, \overline{p})$. Since MLRP implies FOSD, the biases then follow from the fact that the distribution F_C first-order dominates F_L and is first-order dominated by F_H.■

The intuition for the preceding proposition is as follows. In state H (L), agents assigned the treatment draw their direct physiological effects from a distribution that dominates (is dominated by) the distribution from which the control group draws. In expectation, this causes them to assess a higher (lower) probability of state H and to have more (less) favorable brain-modulated physiological responses.

The next proposition presents general conditions under which there will be upward bias in both states.

Proposition 3 *If $\frac{f_L}{f_C}$ and $\frac{f_H}{f_L}$ are both strictly increasing (MLRP), then the expected treatment-control health quality difference is greater than the expectation of the direct physiological effect in both states L and H.*

Proof. From FOSD it follows that in order to establish the entire result we must simply verify β is strictly increasing on $(\underline{p}, \overline{p})$. Rearranging terms in equation (8) we find that β is strictly increasing iff

$$\begin{aligned} R'_H(p)\left[tR_L(p)+(1-t)\right] - R'_L(p)\left[tR_H(p)+(1-t)\right] &> 0 \\ t\left[R_L(p)R'_H(p) - R'_L(p)R_H(p)\right] + (1-t)\left[R'_H(p) - R'_L(p)\right] &> 0 \\ tR_L(p)^2\left[\frac{f_H(p)}{f_L(p)}\right]' + (1-t)\left[R'_H(p) - R'_L(p)\right] &> 0 \end{aligned}$$

From the second MLRP condition we know the first term in the preceding equation is positive. Considering the second term, note that it too is positive under the two MLRP conditions since:

$$R_H(p) - R_L(p) = \frac{f_H(p)}{f_C(p)} - \frac{f_L(p)}{f_C(p)} = \frac{f_H(p) - f_L(p)}{f_C(p)} = \frac{\frac{f_H(p)}{f_L(p)} - 1}{f_C(p)/f_L(p)}. \blacksquare$$

The intuition for the preceding proposition is as follows. In both states, agents assigned the treatment draw their direct physiological effects from a distribution that dominates the distribution from which the control group draws. In expectation, this causes them to assess a higher probability of state H and to have more favorable brain-modulated physiological responses.

Finally, we describe a case in which $\mu_L = \mu_C$, so that in state L the treatment has no positive effect on health relative to the control. Yet, under the stated assumptions, the expected treatment-control health quality difference is positive in state L. The construction is as follows. Suppose $f_L(p) > f_C(p)$ for $p > 3/4$. Consider then Ψ having a convex kink at zero at the point $\beta(3/4)$. Under this type of convex brain-modulated physiological response, the treatment has a stronger mental effect than the control even in state L.

Remark 1 *Let* $\mathcal{P} = [0,1]$, *with* $f_H = 2p$; $f_L = 1$; $f_C = 4p$ *for* $p \leq 1/2$; *and* $f_C = 4(1-p)$ *for* $p > 1/2$ *implying* $\mu_L = \mu_C$. *Suppose further*

$$\begin{aligned} \Psi(X) &\equiv \max\{X - X^*, 0\}, \\ X^* &\equiv \frac{\pi_L}{2} + \left[\frac{\lambda + \frac{\lambda t}{2}}{1 + \frac{\lambda t}{2}}\right]\left[\frac{2\pi_H}{3} - \frac{\pi_L}{2}\right]. \end{aligned}$$

Then the expectation of the treatment-control health quality difference is greater than the direct physiological effect in state L, as well as H.

Proof. We first verify β is increasing, implying so too is M. Let

$$\begin{aligned}\gamma(p) &\equiv \frac{tR_H(p)+(1-t)}{tR_L(p)+(1-t)} \;\forall\; p \in (\underline{p},\overline{p}) \\ &\Rightarrow \ln[\gamma(p)] = \ln[tR_H(p)+(1-t)] - \ln[tR_L(p)+(1-t)] \\ &\Rightarrow \frac{\gamma'(p)}{\gamma(p)} = \frac{[tR_H(p)+(1-t)]'}{[tR_H(p)+(1-t)]} - \frac{[tR_L(p)+(1-t)]'}{[tR_L(p)+(1-t)]}.\end{aligned}$$

From equation (8), it follows the slope of β is equal to the slope of γ. Consider first $p < 1/2$. We have the following increasing function:

$$\gamma(p) = \frac{tR_H(p)+(1-t)}{tR_L(p)+(1-t)} = \frac{1-\frac{t}{2}}{\frac{t}{2p}+(1-t)}$$

Consider next $p > 1/2$. We have the following increasing function:

$$\gamma(p) = \frac{tR_H(p)+(1-t)}{tR_L(p)+(1-t)} = \frac{2tp+4(1-t)(1-p)}{t+4\,(1-t)\,(1-p)}.$$

The rest of the proof follows from the description preceding the remark.■

4 Concluding Remarks

We illustrate a fragility associated with double-blind medical RCTs, often viewed as the gold standard in medicine for estimating non-placebo direct physiological effects (characteristic or specific effects). We show that when positive expectancy about future health quality leads to better present-day health quality, e.g. due to stress reduction, then the expectation of mental effects cannot be presumed equal across treatment and control groups in RCTs, since beliefs will vary systematically with the distribution of direct physiological effects. As was shown, even if a treatment has zero expected effect, the fact that its probability distribution differs from that of the control can result in positive expected treatment-control differences.

The expectancy of medical subjects is related to their assessment of the probability of approval and production of a drug, captured by the model parameters (π_L, π_H). These parameters are likely

to vary with the financial constraints of companies, regulation, and governmental funding capacity. It follows that differences in experimental outcomes are to be expected even if subjects across countries are physically and emotionally identical. That is, optimism and hope, will rationally vary with the institutional setting even with identical subject characteristics.

Finally, since rational expectancy capitalizes anticipations of future improvements in a drug therapy, it might be difficult for next-generation drugs to prove better physiological effects if their pronounced placebo effect on prior-generation drugs is not taken into account.

References

[1] Angrist, Joshua D. and Jorn-Steffen Pischke, 2009, *Mostly Harmless Econometrics: An Empiricist's Companion,* Princeton University Press.

[2] Deaton, Angus, 2010, Instruments, Randomization, and Learning about Development, *Journal of Economic Literature* 48 (2), 424-455.

[3] Grünbaum, Adolf, 1986, The Placebo Concept in Medicine and Psychiatry, *Psychological Medicine* 16 (1), 19-38.

[4] Haygarth, John, 1801, *Of the Imagination as a Cause and as a Cure of Disorders of the Body: Exemplified by Fictitious Tractors and Epidemical Convulsions.* Bath: Crutwell.

[5] International Conference of Harmonization, 2000, Choice of Control Group and Related Issues in Clinical Trials E10, Department of Health and Human Services: Center for Biological Evaluation and Research.

[6] Kaptchuk, Ted J., 1998, Powerful Placebo: The Dark Side of the Randomised Controlled Trial, *The Lancet* (351), 1722-1725.

[7] Malani, Anup, 2006, Identifying Placebo Effects with Data from Clinical Trials, *Journal of Political Economy* 114 (2), 236-256.

[8] Pincus, Theodore, 1998, "Rheumatoid Arthritis: Disappointing Long-Term Outcomes Despite Successful Short-Term Clinical Trials", *Journal of Clinical Epidemiology*, 41, 1037-1041.

[9] Rothwell PM, 2005, "External Validity of Randomised Controlled Trials: To Whom do the Results of this Trial Apply?", *Lancet*, 365, 82-95.

[10] Rothwell PM, 2005, "Factors That Can Affect the External Validity of Randomised Controlled Trials", *PLOS Clinical Trials*, 1-5.

[11] Stewart-Williams, Steve, and John Podd, 2004, The Placebo Effect: Dissolving the Expectancy versus Conditioning Debate, *Psychological Bulletin* 130 (2), 324-340.