

Dynamic Treatment Choices and Selection into RCTs: Evidence from Prostate Cancer *

Andrew Capron Barton H. Hamilton

Brian McManus Juan Pantano

Justin G. Trogdon

October 16, 2026

Abstract

We study how voluntary participation and treatment protocols affect the interpretation of randomized controlled trials (RCTs) comparing existing medical treatments. When patients can obtain their preferred treatment outside the trial, those closer to indifference between treatment strategies are more likely to participate. When treatment is dynamic, protocols governing when patients may switch from their initially assigned treatment affect both enrollment and potential outcomes. We develop and estimate a dynamic model of trial participation and treatment choices using aggregate data on volunteers and eligible non-volunteers from a prostate cancer study. Selection operates primarily through patients' preferences over treatment side effects rather than differences in treatment effectiveness. Under the observed protocol, estimated volunteer and population treatment effects are similar for 10-year prostate cancer mortality but differ for metastasis and 15-year mortality. Alternative switching protocols change both participation and treatment effects.

*Capron: Department of Economics, University of North Carolina at Chapel Hill, acapron@unc.edu; Hamilton: Olin Business School, Washington University in St. Louis, hamiltonb@wustl.edu; McManus: Department of Economics, University of North Carolina at Chapel Hill, mcmanusb@unc.edu; Pantano: The University of Hong Kong, jpanta@hku.hk; Trogdon: Department of Health Policy and Management, Gillings School of Global Public Health, University of North Carolina-Chapel Hill, trogdonj@email.unc.edu. We thank Michael Darden, Hanming Fang, Gautam Gowrisankaran, Mike Keane, Qi Li, David Meltzer, Nick Papageorge, Jin Zhou, for helpful comments and seminar participants at UChicago Med, HKU Med, Johns Hopkins University, Indiana University, Lehigh University, Arizona State University, University of Arizona, City University of Hong Kong, Universidad Torcuato Di Tella and UNC-Chapel Hill as well as conference participants at the Health Econometrics Workshop (HEW), the American Society of Health Economics (ASHEcon), the Asian Workshop on Econometrics and Health Economics (AWEHE), the South East Health Economics Study Group (SHESG), the North American Summer Meetings of the Econometric Society (NASMES) at UCLA and the Dynamic Structural Econometrics Conference at HKU. This study was approved by the institutional review board of the University of North Carolina at Chapel Hill.

Keywords: Selection into RCTs, prostate cancer, dynamic decisions in healthcare
JEL Codes: C90, I12

1 Introduction

Randomized controlled trials (RCTs) are traditionally viewed as the gold standard for treatment evaluation. By randomly assigning interventions, RCTs provide strong internal validity, allowing researchers to estimate treatment effects without confounding from unobserved characteristics of enrolled participants. In healthcare, RCTs underlie the evidence-based approach to medicine, in which treatment effects estimated from trial participants are used to inform clinical practice. However, recent research has questioned whether RCT results can be generalized to broader populations when treatment effects are heterogeneous and trial participants are not representative of the patients who may receive the treatment (Heckman, 1992; Basu et al., 2007; Malani, 2008).

We develop a framework that combines patient behavior and outcomes from RCT participants and non-participants to recover average treatment effects for the full eligible patient population. Our starting point is that participation in RCTs is voluntary due to human subjects concerns, creating the possibility that patients who accept randomization differ systematically from those who decline. Consequently, an RCT identifies what we term the volunteer average treatment effect (VATE), which may differ from the average treatment effect (ATE) in the eligible population when treatment effects are heterogeneous. We address this selection problem by modeling patients' decisions to participate in the RCT and their treatment choices outside the trial.

We apply our framework to comparative-effectiveness RCTs, which evaluate treatments that are already available to patients outside the trial.¹ This setting generates a distinctive form of selection into randomization. In trials of novel treatments, patients may volunteer

¹Examples of well known comparative effectiveness RCTs test local vs. general anaesthesia for lower-extremity surgery (Johnson et al., 2016), different doses of existing drugs (e.g., statins for primary prevention (Yebo et al., 2019)), investigate off-label use of approved drugs (e.g., antipsychotics for dementia (Yunusa et al., 2019) or schizophrenia (Huhn et al., 2019)), drugs for prevention of migraine headaches (Jackson et al., 2015)), or compare behavioral interventions (e.g., meditation for diet and weight loss (Katterman et al., 2014)), non-pharmacological treatment of depression (Farah et al., 2016), Alcoholics Anonymous for alcohol use disorder (Kelly et al., 2020), Weight Watchers for weight loss (Thomas et al., 2017)).

to gain access to a treatment that would otherwise be unavailable, creating selection on potential gains from participation.² In comparative-effectiveness trials, by contrast, patients can obtain their preferred treatment outside the RCT. Patients with strong preferences for one treatment therefore have relatively little incentive to accept randomization, while those who are closer to indifferent between treatments are more willing to participate. Indeed, the protocol for the trial we study explicitly states: “Men should not undergo randomisation unless they are able to view all treatments as reasonably equivalent” (Hamdy et al., 2016).³ The resulting selection on indifference can make RCT volunteers systematically different from the patients whose treatment decisions the trial is intended to inform.

Previous research has recognized the potential for voluntary participation and heterogeneous treatment effects to limit the external validity of RCTs (Heckman (1992); Basu et al. (2007); Malani (2008); Sianesi (2017)). Proposed solutions include using instrumental variables to recover marginal treatment effects and other population parameters (Heckman and Vytlacil (2001), Heckman and Vytlacil (2005); Basu et al. (2007)), combining experimental and observational data (Athey et al. (2020)), and using variation in experimental participation incentives or other design features (List (2026)). Our approach builds on this literature by jointly modeling selection into the RCT, treatment choices outside the RCT, and subsequent treatment decisions when treatments are dynamic. This approach is particularly useful when the treatment strategies available within the RCT differ from those pursued outside it, making it necessary to account for both endogenous participation and differences in the treatment choices permitted under alternative experimental protocols.

Our framework allows patients to differ in their health characteristics, preferences over treatment side effects, and match with available treatments. Patients decide whether to accept randomization into an RCT or remain outside the trial, where they can select their

²The second major class of RCT investigates the effectiveness of a new treatment (e.g., a COVID-19 vaccine). In these RCTs, the treatment in at least one arm of the trial is only available to trial participants. Participants may volunteer in the hope of receiving the experimental drug, implying a preference for the experimental arm of the RCT rather than indifference.

³We assume that 1) patients and their doctors have a good understanding of the health/medical processes that patients experience under each arm of the RCT and that 2) patients and doctors have rational expectations about sources of uncertainty in treatment. Together, these features of comparative effectiveness RCTs imply that volunteers for these RCTs must be (close to) indifferent between the trial’s treatment arms, otherwise they would take their preferred treatment outside of the RCT.

preferred treatment strategy. When treatment choices are dynamic, patients also consider the option of switching treatments in response to disease progression.⁴ Within an RCT, a protocol may restrict when these switches are permitted. The participation decision therefore depends on both the prospect of random assignment to an initially less desirable treatment and the restrictions that trial protocols place on subsequent choices. Moreover, the switching protocol can affect potential treatment outcomes, so that the average treatment effect depends on the protocol as well as the population being evaluated. Our framework distinguishes the treatment effect among RCT volunteers under the trial protocol from the corresponding effect in the full eligible population and from effects under alternative protocols, including unrestricted switching. By modeling these decisions jointly, we can use data on RCT volunteers and non-volunteers to evaluate treatment effects for populations and protocols not directly observed in the trial.

In the context of our empirical application to prostate cancer treatment, patients choose between an aggressive treatment strategy (e.g., surgery or radiation) and a minimally aggressive surveillance strategy with the option to switch to more aggressive treatment in the future. Patients who enroll in an RCT are randomly assigned an initial treatment strategy. Under surveillance, patients weigh the benefits of switching to an aggressive treatment that is more likely to cure the disease but also more likely to generate negative side effects. This choice depends on how patients weigh the trade-off between curative effectiveness and side effects, which can vary with unobserved characteristics (preferences and patient match to treatment) as well as observable characteristics such as age and indicators of disease severity. The patient’s willingness to switch ultimately affects the risks of metastasis and cancer mortality. To capture these choices, we specify a dynamic model of patient behavior. Patients outside the RCT choose freely between treatment strategies and, if they begin with surveillance, choose when and whether to switch to aggressive treatment. Patients compare the value of these unconstrained choices with participation in the RCT, which includes randomization and a protocol that may restrict when patients can switch treatments.

⁴[Chan and Hamilton \(2006\)](#) also develop a dynamic model of patient decisions within an RCT, emphasizing learning about treatment effectiveness and privately observed side effects as determinants of experimental attrition. Unlike their analysis, our framework focuses on selection into the RCT and the effects of treatment-switching protocols on participation and outcomes.

The empirical model links observed RCT participation, initial treatment choices, and the timing of treatment switching to patients’ mortality and metastasis outcomes. These moments, together with calibrated inputs on the utility effects of treatment side effects, inform the model’s parameters describing disease progression, treatment effectiveness, and the distributions of patients’ preferences and treatment-match values. With the estimated parameters, we can simulate treatment choices and potential outcomes for volunteers and non-volunteers under the same treatment protocol, as well as evaluate outcomes under alternative protocols.

We apply our method to published RCT and observational data on prostate cancer treatment. Our main data source is the Prostate Testing for Cancer and Treatment (ProtecT) study, which followed about 2,600 newly diagnosed UK prostate cancer patients eligible for the study. The centerpiece of ProtecT is an RCT that randomized about 1,600 men across surveillance and active treatment arms ([Hamdy et al., 2016](#)). In the RCT, some patients assigned to surveillance subsequently receive active treatment as their disease progresses. The key outcomes are mortality and metastasis related to the patient’s prostate cancer. In addition to the men randomized across treatments in the RCT, approximately 1,000 eligible men declined randomization and instead selected their own treatment strategies. These men comprise the ProtecT Choice sample, our main source of observational data. Because Choice participants faced the same enrollment decision as RCT volunteers, their treatment choices and outcomes provide particularly useful information about selection into the trial and differences between patients who accept and decline randomization.

We estimate the empirical model by matching observed moments on RCT participation, treatment choices, and health outcomes to simulated patient histories. The estimated model closely matches the share of eligible patients who participate in the RCT (62.0% in the data and 62.3% in the model), as well as important patterns in cancer mortality and metastasis. The model captures treatment switching relatively well, especially in the early years following diagnosis. Our estimates reveal substantial selection on patients’ aversion to treatment side effects. Patients who decline randomization and choose immediate treatment tend to have relatively low aversion to side effects, while those who decline and choose surveillance tend to have high aversion. By contrast, selection on patients’ match with treatment is relatively

limited, and volunteers and non-volunteers have similar observed characteristics such as age and initial disease severity. Thus, the model identifies unobserved preferences over treatment side effects as the key source of selection into the RCT.

Using the estimated model, we compare the volunteer average treatment effect (VATE) to the corresponding average treatment effect (ATE) for the eligible patient population under the same RCT protocol. We focus on the effects of surveillance relative to immediate treatment on prostate-cancer mortality and metastasis. Under the ProtecT protocol, the estimated VATE and ATE for 10-year prostate-cancer mortality are approximately equal. There are differences, however, across other outcomes and time horizons. At 15 years, the estimated mortality effect for the full eligible population exceeds the effect among RCT volunteers. Similarly, the VATE understates the population effect on metastasis at both 10 and 15 years. These results are driven by the strong distaste for treatment side effects by many individuals who opt-out of treatment, generating greater rates of disease progression and illness than among RCT volunteers, who are more likely to switch to active treatment when permitted by the experimental protocol.

We also use the estimated model to evaluate counterfactual RCT designs that vary the severity threshold at which patients initially assigned to surveillance may switch to active treatment. These changes affect both the treatment outcomes that would be observed in the RCT and the composition of patients willing to participate. Patients with relatively low aversion to treatment side effects become more willing to enroll when the protocol permits earlier switches to active treatment. In contrast, patients with high aversion to side effects remain relatively unlikely to volunteer, even under less restrictive protocols, because they do not value the option to switch earlier and still prefer to avoid potential randomization to immediate active treatment. Thus, relaxing switching restrictions can increase participation among some patient types without eliminating selection into the RCT. The counterfactual results also show that differences between the ATE and VATE depend on the protocol, with particularly notable differences for longer-term cancer mortality and metastasis. These findings highlight the importance of considering how RCT design affects both the potential outcomes under the treatments being compared and the composition of the experimental sample.

2 From VATE to ATE

2.1 Setting

Consider a population of individuals who have just been diagnosed with a disease. The framework is general and can apply to either lethal or non-lethal diseases, mild or serious. Diagnosis occurs between ages $\underline{a}$ and $\bar{a}$; we use a_0 to denote a patient's age at diagnosis. Time is discrete. After diagnosis, the patient makes choices about its treatment. There are two treatment options available: advanced treatment (A), which involves relatively invasive medical intervention, and basic treatment (B), which is less invasive and, in some cases, may be limited to monitoring. Let Y_i denote a health outcome of interest, for example 10-year survival. We define $Y_{i,A}$ and $Y_{i,B}$ as the corresponding *potential* outcomes associated with, respectively, treatment choices A and B . Because individual treatment choices are endogenous, observational analysis of treatment choices and outcomes is in general insufficient to identify treatment effects. Indeed, it is typically necessary to conduct an RCT to convincingly estimate the effectiveness of A vs B in the diagnosed population.

In a comparative effectiveness trial, the estimand of interest is typically the Average Treatment Effect of treatment B relative to treatment A , on outcome Y , for the full population of newly diagnosed patients: $ATE = E[Y_{i,B} - Y_{i,A}]$. That is the effect on outcome Y of treating everyone with B instead of A . In general though, the RCT identifies only the Volunteer Average Treatment Effect (VATE) $= E[Y_{i,B} - Y_{i,A} | S = 1]$, where S is an indicator for voluntary enrollment into the RCT that equals one for those who accept the invitation to join the RCT and zero for those who decline it.

At diagnosis, patients are characterized by their age, a_0 ; the severity of their disease, ω_{a_0} ; their preferences, α , over treatment side effects that they could experience under treatment A and/or B ; and their match value to treatments A and/or B , $\varphi = (\varphi_A, \varphi_B)$. The α and φ are the key sources of patient unobserved heterogeneity in the model. They are observed by the patient and his doctor but are not necessarily available in the RCT or observational data. Let $f(a_0, \omega_{a_0}, \alpha, \varphi)$ represent the joint density of these characteristics in the population of newly diagnosed patients.

In this paper, we focus on two aspects of treatment effects and experimental design.

First, we study the estimation challenges that arise from the patients' characteristics (α, φ) , which we treat as unobservable to the researcher. While the ATE provides the treatment effect over the full population's distribution of (α, φ) , VATE only integrates it over the *conditional* distribution of these unobservables among those who join the RCT. Differences in patients' observable characteristics, (a_0, ω_{a_0}) , which can also affect RCT participation, are more straightforward to address through approaches such as reweighting volunteers' characteristics to resemble the full population.

Second, we focus on treatments where A , B , or both are dynamic, in the sense that duration in a treatment may matter for outcomes, and patients may have the opportunity to switch between treatments after treatment begins. We return briefly to the case of static treatments below. When at least one dynamic treatment is involved, RCTs typically will establish a protocol $\underline{\omega}$ that specifies if and when a patient is allowed to switch into another treatment given the initial treatment assignment randomized in the RCT. We use the protocol notation $\underline{\omega}$ to suggest a connection between disease progression and patient opportunities to switch treatments, but other types of protocols are also possible, e.g. duration or a number of doses with the initial treatment assignment. If this protocol was not in place, the RCT would be less informative as patients could immediately switch to their preferred treatment, re-instating the endogeneity in treatment choices prevailing outside the RCT. Therefore, when treatments are dynamic note that we must take into account the protocol $\underline{\omega}$.

In general then, potential outcomes under some or all of the treatments could be affected by this protocol and we denote them $(Y_{i,B}(\underline{\omega}), Y_{i,A}(\underline{\omega}))$. Therefore, the estimands must be indexed by the protocol as well. The RCT identifies a VATE under that specific protocol $\underline{\omega}$: $\text{VATE}(\underline{\omega}) = E[Y_{i,B}(\underline{\omega}) - Y_{i,A}(\underline{\omega}) | S = 1]$. This could be different from VATEs under different protocols $\underline{\omega}'$, or the VATE in the absence of a protocol, $\underline{\omega} = 0$. Similarly, one must then consider $\text{ATE}(\underline{\omega}) = E[Y_{i,B}(\underline{\omega}) - Y_{i,A}(\underline{\omega})]$. But the protocol $\underline{\omega}$ implemented in the RCT may or may not be the one of policy interest. Therefore, one must not only overcome the selected population problem but also the fact that the potential outcomes are themselves affected by an RCT switching protocol that might not be the same as envisioned in the target policy environment. For example, researchers may want to set a relatively restrictive $\underline{\omega}$ to achieve some informational objectives in the RCT, while practitioners may consider less restrictive

treatment strategies for patients with additional concerns (e.g., risk attitudes) beyond the RCT's scope.

$\text{ATE}(\underline{\omega})$ is a weighted average of the average treatment effect among volunteers $\text{VATE}(\underline{\omega})$ and the one among non-volunteers, $\text{nVATE}(\underline{\omega})$, weighted by the share of volunteers in the diagnosed population $\Pr(S = 1|\underline{\omega})$:

$$\text{ATE}(\underline{\omega}) = \Pr(S = 1|\underline{\omega})\text{VATE}(\underline{\omega}) + (1 - \Pr(S = 1|\underline{\omega}))\text{nVATE}(\underline{\omega})$$

While $\Pr(S = 1|\underline{\omega})$ is easily identified by the patient population response to the RCT invitation, and $\text{VATE}(\underline{\omega})$ is directly identified by the RCT, $\text{nVATE}(\underline{\omega})$ is not identified, absent a model of selection into RCTs. This is the case even if one has access to data on treatment choices and outcomes among non-volunteers who declined the invitation to join the RCT. Again, there are typically two problems: first, the data from non-volunteers involves endogenous treatment choices, as these were freely decided by patients, rather than randomized. Second, regardless of whether the ultimate estimand of interest is $\text{ATE}(0)$ or $\text{ATE}(\underline{\omega})$ for $\underline{\omega} > 0$, the (endogenous) choices freely made outside an RCT are typically not restricted by RCT-like protocols such as $\underline{\omega}$. Therefore, at best, they could only identify $\text{nVATE}(0)$ rather than the needed $\text{nVATE}(\underline{\omega})$.

Given these identification problems, a model of selection into RCTs is necessary to characterize how selection on (α, φ) generates differences between the VATE and ATE and, in particular, to recover treatment effects for patients who do not volunteer for the RCT. When treatment choices are dynamic, nesting this selection decision within a dynamic model provides additional structure for understanding the sources of selection. In particular, the model allows us to distinguish selection arising from patients' preferences over initial treatment assignments from selection arising from the possibility that the RCT protocol will constrain future treatment switching. The dynamic structure also allows us to decompose differences in potential outcomes between patients inside and outside the RCT into differences due to the composition of patients who select into the RCT and differences generated by the treatment paths available under the RCT protocol.

Using data on treatment choices and outcomes by RCT volunteers and non-volunteers

it is possible to identify a structural model that can serve as the basis for recovering the ATE. The model features selection into the RCT, treatment choice outside of the RCT and, for dynamic treatments with switching option, also the treatment choices in the RCT after initial randomization, if and when the switching protocol conditions are eventually satisfied by the patient. With that estimated model in hand, one can then identify the $\text{nVATE}(\underline{\omega})$ and thus $\text{ATE}(\underline{\omega})$. The model can then be used to evaluate the relevant non-volunteer population through the two alternative treatment policy functions and estimated disease evolution process using the model. Similarly, one can directly use the model to obtain estimates of other, perhaps more policy-relevant estimands such as $\text{ATE}(\underline{\omega}')$ for $\underline{\omega}' \neq \underline{\omega}$ or $\text{ATE}(0)$.

2.2 RCT enrollment decision

Newly diagnosed patients are offered the opportunity to enroll in an RCT (“**rct**”) that will test the health effects of initial assignment to B versus A . In the RCT, patients who accept the invitation and join the trial are randomized to A with known probability p_A and are assigned to treatment arm B with probability $p_B = 1 - p_A$. We summarize the patient’s initial observed state variables at age a_0 with the vector $\chi_{a_0} = \{a_0, \omega_{a_0}\}$. This initial state vector contains the patient’s age at diagnosis (a_0) and the severity of the disease at the time of diagnosis (ω_{a_0}).

To decide whether to enroll in the RCT, newly diagnosed patients in the population $f(\chi_{a_0}, \alpha, \varphi)$ compare the value of accepting and declining RCT participation. We assume that, in addition to the possible values under assignment to A or B given RCT participation captured in $V_A(\chi_{a_0}, \mathbf{rct})$ and $V_B(\chi_{a_0}, \mathbf{rct})$, each enrolled patient receives an additional utility shock, ϵ . This shock represents an additional benefit that the patient receives from interacting with medical staff at the time of enrollment, for example due to a patient’s feeling of obligation to help his physician recruit patients into the RCT.⁵ Each patient draws a value of $\epsilon \in \mathcal{R}$ from f_ϵ and observes its value prior to the RCT enrollment choice. In the

⁵The patient receives no extra medical attention or care within the RCT relative to outside of it. If so, this could affect a patients beliefs about how ω is monitored or whether mortality risks are different within the RCT.

next subsections we discuss in more detail the value functions (V_A, V_B) both in and out of the RCT, which we take as given here.

After observing his realization of ϵ , the patient can compute the value of enrolling in the RCT with protocol $\underline{\omega}$:

$$V(\chi_{a_0}, \epsilon, \mathbf{rct}, \underline{\omega}) = p_A V_A(\chi_{a_0}, \mathbf{rct}, \underline{\omega}) + (1 - p_A) V_B(\chi_{a_0}, \mathbf{rct}, \underline{\omega}) + \epsilon. \quad (1)$$

The value of declining the RCT offer, $V(\chi_{a_0}; \mathbf{out})$, is the value of remaining free to choose the best course of action at any time, avoiding randomization of initial treatment and when dynamic treatments are involved, unconstrained by any decision protocol in any subsequent decision-making period.

$$V(\chi_{a_0}; \mathbf{out}) = \max\{V_A(\chi_{a_0}; \mathbf{out}), V_B(\chi_{a_0}; \mathbf{out})\} \quad (2)$$

A patient will then enroll in the RCT if and only if the value of joining the RCT exceeds the value of staying out of it:

$$S = 1 \iff V(\chi_{a_0}, \epsilon; \mathbf{rct}, \underline{\omega}) \geq V(\chi_{a_0}; \mathbf{out}). \quad (3)$$

To better understand the selection issues in enrolling patients in the RCT, consider a comparison of $V(\chi_{a_0}, \epsilon; \mathbf{rct}, \underline{\omega})$ and $V(\chi_{a_0}; \mathbf{out})$ for a patient who draws $\epsilon = 0$. For such a patient, enrollment is at most equally worthwhile to declining the invitation and only if he is indifferent between treatments A and B in period a_0 . Indeed, if the patient was not indifferent, he would prefer to choose A or B and receive the associated value with certainty rather than facing randomization between the preferred option and the alternative. This indifference between A and B implies that enrolled patients must have a knife's edge combination of $(a_0, \omega_0, \alpha, \varphi)$ values, and therefore, are not representative of the full patient population.

In addition to indifference between A and B outside of the RCT, a patient with $\epsilon = 0$ willing to enroll in the RCT must have $V_d(\chi_{a_0}; \mathbf{rct}, \underline{\omega}) = V_d(\chi_{a_0}; \mathbf{out})$ for any treatment $d \in \{A, B\}$ for which the RCT specifies a switching protocol $\underline{\omega}$ to switch treatment away

from d . This condition holds when the RCT’s protocol restrictions on switching have no impact on preferred treatment choices at any future value of ω_a . If the condition fails to hold, then even a patient with $V_A(\chi_{a_0}, \text{out}) = V_B(\chi_{a_0}, \text{out})$ will prefer either non-RCT option to randomization between RCT treatment arms. Indeed, a patient may not exist that can satisfy all of the conditions for participation in the RCT when $\epsilon = 0$. By including the utility shock ϵ in our model, we ensure that some patients, with sufficiently large ϵ and $V_A(\chi_{a_0}, \text{out})$ close to $V_B(\chi_{a_0}, \text{out})$, are willing to participate in the RCT while preserving the key selection issues that affect voluntary RCT enrollment.

2.3 Dynamic Treatment Choices Inside and Outside the RCT

When treatment is dynamic, patients can switch treatments sequentially so they must consider treatment sequences. Among these settings, it is possible that one of the treatment choices, say “A”, might be absorbing or static, in that no further treatment choices are possible after making treatment choice A. Unless initially choosing or being assigned to A, treatment choice becomes an optimal stopping problem, with the patient choosing when (if ever) to switch from B to A . For yet other diseases, the patient may sequentially combine treatments A and B in spells of various lengths as he or she sees fit.

When dynamics are involved, we denote a patient’s treatment choice at any age a by d_a . A patient makes his first treatment decision during the period of diagnosis, a_0 . The disease’s state evolves and is still captured by the biomarker, ω_a , that measures disease progression. A patient’s value of ω now evolves according to a *controlled* stochastic process $f_\omega(\omega_{a+1}|\omega_a, d_a; \varphi)$.

Patients diagnosed at age a_0 can die at any age between $a_0 + 1$ and $\bar{a}$, whether from the disease of interest or from other causes. If a patient survives to $\bar{a}$, he dies during the next period with probability one. Let H_a be a survival indicator for an individual patient at age a , with $H_a = 1$ if the patient is alive and zero otherwise. Despite its dynamics, the model remains one focused on treatment choice, so the model begins at age a_0 and we abstract away from the periods leading up to diagnosis.

In this dynamic case, we summarize the patient’s state variables at age a with the vector $\chi_a = \{a, \omega_a, \omega_{a_A}, d_{a-1}, a_A, H_a\}$. The state vector contains the patient’s age (a), the current

severity of the disease (ω_a), the patient's treatment choice during the previous period (d_{a-1}), the age at which the current treatment was taken for those who ever took active treatment (a_A if on treatment A or a_B if on treatment B), the disease's severity when the patient chose the current treatment (ω_{a_A} or ω_{a_B}), and the survival indicator (H_a). Patient age increases by one unit for each period the patient survives, and ω_a evolves according to $f_\omega(\omega_{a+1}|\chi_a, d_a, \varphi)$. Survival from a to $a+1$ depends on the avoidance of mortality due to the disease (for potentially lethal diseases) and due to other causes; the survival probability is $f_H(H_{a+1}|a, \omega_{a+1}, d_a, a_{da}, \omega_{a_{da}}, \varphi)$. We include φ in f_H to highlight the role of treatment effectiveness heterogeneity in determining disease-related mortality.

Each period that the patient is alive, he receives flow utility equal to $u(d_a, \chi_a, \alpha) \geq 0$. The patient receives flow utility equal to zero when he is not alive. A patient with a greater value of α has a greater distaste from treatment side effects, which reduce the flow utility from being alive. Treatment side effects could be temporary or permanent and occur after either one or both treatments. Outside of an RCT ("out"), a patient chooses a sequence of d_a values to maximize the expected value of lifetime utility:

$$V(\chi_a, \text{out}) = \max_{\{d_a(\chi_a)\}_{a=\bar{a}_0}^{\bar{a}}} \left\{ E \left[\sum_{a=a_0}^{\bar{a}} \delta^{a-a_0} u(d_a, \chi_a, \alpha) \middle| \chi_{a_0} \right] \right\}. \quad (4)$$

$E[\cdot]$ is with respect to the transitions in the state variables, including the survival transition. δ discounts future periods' utility.

The recursive, dynamic programming representation of the patient's problem outside the RCT is:

$$V(\chi_a, \text{out}) = \max_d \left\{ u(d, \chi_a, \alpha) + \delta E[V(\chi_{a+1}, \text{out}) | \chi_a, d_a = d] \right\} \quad (5)$$

The alternative-specific value functions for $d \in \{A, B\}$ are:

$$V_d(\chi_a, \text{out}) = u(d, \chi_a, \alpha) + \delta E[V(\chi_{a+1}, \text{out}) | \chi_a, d_a = d] \quad (6)$$

The patient's value from choosing between A and B is:

$$V(\chi_a, \text{out}) = \max_d \left\{ V_d(\chi_a, \text{out}) \right\}.$$

Next we discuss a patient's choices, utility and value function within the RCT. In the RCT the initial treatment decision at the time of diagnosis a_0 is randomized. In cases where dynamic treatment choices are possible, a patient's ability to switch away from the initially assigned treatment during periods after a_0 , will typically be constrained by an RCT protocol. Certain conditions must be met for the patients assigned to either A or B to consider switching to the alternative treatment (B or A). Thus, behavior in the RCT is a constrained version of the patient's problem outside the RCT. The RCT in general specifies threshold conditions, $\underline{\omega}$, that the patient's history of biomarker values, $\boldsymbol{\omega}_a = (\omega_{a_0}, \omega_{a_0+1}, \dots, \omega_a)$, must satisfy in order for the patient to consider switching treatment. Let $1_{\underline{\omega}}\{\boldsymbol{\omega}_a\}$ be an indicator function that reports whether a patient with $\boldsymbol{\omega}_a$ satisfies the threshold condition. A simple example of $1_{\underline{\omega}}\{\cdot\}$ would be a comparison of the current ω_a to a scalar threshold. We write $\bar{p}_{\omega} = \Pr(1_{\underline{\omega}}\{\boldsymbol{\omega}_{a+1}\}|\boldsymbol{\omega}_a)$ as the probability that the disease marker evolves between a and $a+1$ in a way that satisfies the RCT's protocol condition(s) for unconstrained decision-making during time period $a+1$.

A patient enrolled in the RCT who is assigned to treatment d receives that treatment in period a_0 . As outside the RCT, the patient immediately experiences any side effects associated with treatment d . If d is absorbing, the patient makes no further treatment choices, and the value of the assignment is the expected present discounted value of remaining lifetime utility under that treatment.

If the patient is assigned to a non-absorbing treatment d at age a_0 , starting in period a_0+1 his value solves the following constrained dynamic optimization problem:

$$V_d(\chi_{a_0+1}, \mathbf{rct}; \underline{\omega}) = \max_{\{d_a(\chi_a)\}_{a=a_0+1}^{a=\bar{a}}} \left\{ E \left[\sum_{a=a_0+1}^{\bar{a}} \delta^{a-a_0} u(d_a, \chi_a, \alpha) \middle| \chi_{a_0+1}, d_{a_0} = d \right] \right\}$$

subject to $d_a(\chi_a) = d$ if $1_{\underline{\omega}}\{\boldsymbol{\omega}_a\} = 0$.

If the patient's ω history does not satisfy the threshold condition ($1_{\underline{\omega}}\{\boldsymbol{\omega}_a\} = 0$), he must remain in treatment d during the current period. If $1_{\underline{\omega}}\{\boldsymbol{\omega}_a\} = 1$, he may choose between switching or continuing with the current treatment. In recursive form, at any age a , the value

of continuing in non-absorbing treatment d in the RCT (either by choice or by protocol) is:

$$\begin{aligned}
V_d(\chi_a, \mathbf{rct}; \underline{\omega}) &= u(d, \chi_a, \alpha) \\
&+ \delta \left\{ \bar{p}_\omega E[\max_{d'} \{V_{d'}(\chi_{a+1}, \mathbf{rct}; \underline{\omega})\} | \chi_a, d_a = d, 1_{\underline{\omega}}\{\omega_{a+1}\} = 1] \right. \\
&\left. + (1 - \bar{p}_\omega) E[V_d(\chi_{a+1}, \mathbf{rct}; \underline{\omega}) | \chi_a, d_a = d, 1_{\underline{\omega}}\{\omega_{a+1}\} = 0] \right\}.
\end{aligned} \tag{7}$$

This expression has three parts. First, the patient receives $u(d, \chi_a, \alpha)$, the flow utility associated with treatment d at age a . Second, in the next period with probability $\bar{p}_\omega$, the patient's disease history ω_{a+1} allows him to make an unconstrained choice between treatments $d' \in (A, B)$. Third, in the next period with probability $(1 - \bar{p}_\omega)$, the patient will be constrained by the RCT protocol to continue with treatment d . The term $V_d(\chi_a, \mathbf{rct}; \underline{\omega})$ also captures the value of being assigned to arm d during period a_0 , when a patient enrolled in the RCT is randomized between treatments A and B .

While this generic framework applies to all comparative effectiveness RCTs with dynamic treatments, in analyzing a particular medical setting, specific choices will need to be made to parameterize the model. In the application below, we adapt this framework to a case study on prostate cancer.

2.4 Static Treatments and Selection into an RCT

In some cases two treatment courses are mutually exclusive in the sense that pursuing either one permanently precludes the other. In this setting treatment choice is effectively modeled as static. When both A and B are static or “single shot” treatments, no choices are made inside the RCT. Patients who choose to volunteer for the RCT simply comply with their randomly assigned treatment and observe how their disease unfolds over time given that set treatment. Given a set of treatment procedures (e.g., dosage levels), the patient's value from assignment to A or B is the expected discounted value of lifetime utility under that treatment. Similarly, the patient outside of the RCT compares the value from each available treatment and chooses the option with the greatest present value.⁶

⁶The treatment procedures are analogous to some aspects of the dynamic protocol $\underline{\omega}$, in that they can capture how different the general treatments A and B are from each other.

When static treatment procedures are appropriately similar within and outside the RCT, researchers have several potential avenues to estimate an ATE. These options follow from the opportunity to specify the static problem as a Generalized Roy Model with an additional selection decision governing RCT participation. Patients’ selection choices may differ in both their potential health outcomes under the two treatments and their preferences over treatment attributes, such as side effects. If microdata are available on RCT participation, randomized treatment assignments and outcomes for volunteers, and treatment choices and outcomes for non-volunteers, one can estimate the structural parameters of the selection model via maximum likelihood and identify the ATE. Alternatively, a method-of-moments estimator can be used when only aggregate statistics are available: the means and variances of outcomes under treatments A and B for volunteers and non-volunteers, the shares choosing A and B outside the RCT, and the share of patients in the population who accept the invitation to participate in the RCT. The treatment choices and outcomes of non-volunteers are therefore informative about the distribution of unobserved heterogeneity (α, φ) that jointly shapes treatment choice and selection into the RCT. This has the practical advantage of allowing researchers to revisit studies for which the underlying microdata are no longer available. Even if most studies do not follow or collect data from those who decline to join the RCT, one can supplement data on volunteers with data from a cohort of diagnosed patients who were not invited to participate in the RCT.

In a recent paper, [List \(2026\)](#) considers related issues of how selection into a static-treatment experiment may provide biased estimates of a treatment effect in a target population. [List \(2026\)](#) discusses several approaches to address the bias, including randomization of participation in the experiment, variation in incentives to participate that can serve as an instrumental variable, and sensitivity analysis potentially based on observable data.⁷ While some of these alternative approaches are appealing in the static case, a structural model becomes particularly useful when treatment practices inside and outside the experiment differ and treatment involves non-trivial dynamic sequences. By modeling selection, treatment

⁷While [List \(2026\)](#) points to [Andrews and Oster \(2019\)](#) as a possible approach based on observables to gauge the importance of selection on unobservables, selection on observables may be uninformative about selection on (α, φ) . If available, one could then model both selection into or out of the RCT and treatment choice, extending the non-parametric Roy model framework of [Heckman and Vytlacil \(2005\)](#) to include an additional selection step.

choice, and subsequent treatment dynamics jointly, the model can be used to evaluate counterfactual treatment procedures that may differ from those observed either inside or outside the experiment.

3 Application: prostate cancer

Prostate cancer is the most common type of cancer among men, with over 330,000 estimated new cases annually in the U.S. and 50,000 new cases annually in the U.K. (Siegel et al., 2026; NHS England Digital, n.d.). Prostate cancer is the second leading cause of cancer death among males, with over 36,000 estimated deaths in the U.S. in 2026 (Siegel et al., 2026). Prostate cancer treatment options can be grouped into watchful waiting strategies (i.e., observation or active surveillance) and treatment. Watchful waiting strategies monitor the progression of the disease; active surveillance uses more tests at more regular intervals than observation alone (i.e., annual digital rectal exams, biannual prostate-specific antigen [PSA] tests, annual prostate biopsies or multiparametric magnetic resonance imaging). Treatment strategies include radical prostatectomy, radiation therapy including external beam radiation and brachytherapy, cryotherapy, androgen deprivation therapy, and chemotherapy. Treatment often leads to serious side effects, including sexual dysfunction, urinary incontinence, and bowel symptoms (Chen et al., 2009; Litwin et al., 1995). Because most men are diagnosed with lower risk prostate cancer and prostate cancer is relatively slow-growing (Epstein et al., 2012), active surveillance can spare men side effects while preserving the option of treatment if the disease progresses.

However, active surveillance comes with a higher risk of metastasis and a slightly higher risk of mortality. The first trial comparing watchful waiting to prostatectomy found lower overall mortality (relative risk [RR] = 0.74, 95% confidence interval [CI]: 0.62 - 0.87) and prostate cancer mortality (RR = 0.55, 95% CI: 0.41 - 0.74) in the prostatectomy arm vs the watchful waiting arm at 23 years of follow up (Bill-Axelson et al., 2018). Two more recent trials, including ProtecT, found no differences in overall survival at 10 years between active surveillance and active treatment among men with low- and intermediate-risk prostate cancer (Wilt et al., 2012; Hamdy et al., 2016). Although mortality was similar between

active surveillance and treatment in the ProtecT trial at 15 years, men assigned to active surveillance were twice as likely to develop metastatic disease compared to those assigned to treatment (Neal et al., 2020; Hamdy et al., 2023).

Based on these and other trials, national guidelines have shifted to advocating active surveillance (Loeb et al., 2018; Carroll and Mohler, 2018). Current National Comprehensive Cancer Network (NCCN) treatment guidelines prefer active surveillance for men with low- and very low-risk tumors. Active surveillance is also a first-line option for men with favorable intermediate risk. By 2015, more than 40% of US men with low-risk cancer, and 10% of men with intermediate-risk cancer, initially chose active surveillance (Mahal et al., 2019).

In summary, prostate cancer is amenable to dynamic economic modeling due to its relatively slow growth rate that provides time to assess treatment options; similar expected survival across treatment options that highlights the role of preferences in decision-making; and a clear trade-off between active treatment, which gives more certainty regarding metastasis and survival at the cost of more side effects, and active surveillance, which avoids side effects and preserves the option value of taking aggressive treatment if the cancer worsens.

4 Data

4.1 Disease states, treatment choices, and health outcomes

Our main data were statistics publicly reported from the UK’s ProtecT study (Hamdy et al., 2016; Neal et al., 2020; Hamdy et al., 2023). Between 1999 and 2009, ProtecT recruited men for an RCT studying prostate cancer treatment. Men eligible for ProtecT had PSA levels that did not suggest immediate treatment (i.e., below 20), were between ages 50 and 69 years, had Gleason scores of 6 to 10, and had no spread of the cancer beyond the prostate (i.e., localized cancer).

Of the 2,664 eligible men, ProtecT enrolled 1,643 men who agreed to randomization and were assigned to active surveillance [N=545] or immediate treatment (radical prostatectomy [N=553] or radiotherapy [N=545]). The men in the active surveillance arm had their PSA measured every 3 months in the first year and every 6 to 12 months thereafter. Unlike two

earlier trials (Bill-Axelson et al., 2018; Wilt et al., 2012), ProtecT included a protocol for men in the active surveillance arm to consider transitions to treatment when certain triggers based on changes in PSA levels were met.⁸ After 10 years, 55% of men originally assigned to active surveillance in the trial received some form of treatment (Hamdy et al., 2016).

Adherence to original assignment was incomplete. Within nine months, 482 of the 545 men assigned to active monitoring, 391 of the 553 assigned to prostatectomy, and 405 of the 545 assigned to radiotherapy had begun their assigned protocols. The original assignment counts are useful for understanding the experimental design, but they do not determine the initial treatment frequencies used in our empirical model. Because the model does not separately represent initial non-adherence, we define initial treatment using the treatment-received classifications in Neal et al. (2020). Under this alternative classification, 1,607 randomized men are classified as receiving active monitoring [N=628], prostatectomy [N=488], or radiotherapy [N=491].

We supplemented the ProtecT RCT data with statistics from the ProtecT “Choice” sample, which followed men who met the eligibility criteria but opted out of randomization (N=958).⁹ Men in the ProtecT Choice sample selected active surveillance (N=507), radical prostatectomy (N=262), or radiotherapy (N=189). Choice Sample patients in active surveillance had no restrictions on when they could switch to surgery or radiotherapy. Nevertheless, relative to the RCT, a comparable proportion of patients who initially chose active surveillance switched to treatment in the Choice sample.

We gathered statistics on ProtecT samples from Hamdy et al. (2016), Neal et al. (2020), and Hamdy et al. (2023). Patients who elected randomization (RCT sample) were similar to patients who declined randomization (Choice sample) in terms of age, initial PSA levels, and Gleason scores (Neal et al., 2020, Table 1). This suggests that selection into the RCT based on these observable characteristics was minor.

Regarding treatment choices, we use the treatment-received classifications and subsequent switching statistics reported by Neal et al. (2020). For patients in the RCT, we use the

⁸Specifically, a PSA value of 20+ or an increase of 50% or more in a year triggered a review by the provider team. To keep the state space manageable, we model the trial protocol as an absolute threshold (e.g., when PSA rises above 20).

⁹The remaining 63 eligible men did not participate in the trial or “Choice” sample.

reported numbers at risk and cumulative numbers receiving radical treatment over time. Let t represent years since the second PSA test following diagnosis, which is the starting point used by [Neal et al. \(2020\)](#). For the 628 RCT patients classified as initially receiving active surveillance, they report these statistics at $t \in \{2, 4, 6, 8, 10\}$. After adjusting the reported risk sets for attrition, the cumulative probability of switching to treatment reaches 48% after ten years. This is similar to the path followed by patients who initially received active surveillance in the Choice cohort, for whom the corresponding ten-year probability is 44% ([Neal et al., 2020](#), Figure 1).

For health outcomes, we use the treatment-received incidence rates reported by [Neal et al. \(2020\)](#). Among RCT patients initially receiving active surveillance, the ten-year metastatic-disease rate is 6.0 per 1,000 person-years, compared with 2.7 per 1,000 person-years for those receiving treatment. The corresponding rates are 6.1 and 2.5 per 1,000 person-years in the Choice cohort. Prostate-cancer mortality is also higher under active surveillance: 1.8 versus 0.64 per 1,000 person-years in the RCT cohort and 2.2 versus 0.89 per 1,000 person-years in the Choice cohort ([Neal et al., 2020](#), Table 2). We derive non-cancer mortality by subtracting prostate-cancer mortality from all-cause mortality. Five-year RCT prostate-cancer survival comes from [Hamdy et al. \(2016\)](#), while the 15-year RCT metastasis and mortality statistics come from [Hamdy et al. \(2023\)](#).

In our econometric analysis, we match published aggregate moments from ProtecT trial to corresponding predictions from our empirical model. In Appendix Table 3, we provide a complete list of the ProtecT moments we use in estimation. The values in the ‘Data’ column of Appendix Table 3 come directly from or are derived from the published studies described above. Because the studies report outcomes in different formats and over different horizons, we adjust the published statistics to account for trial attrition, combine surgery and radiotherapy into treatment A , convert annual event rates into cumulative 10- and 15-year risks, and align baseline age and PSA measures with the model’s scale. Appendix Section B provides additional details.

4.2 Calibrated values for patient preferences

We make assumptions on how patients’ per-period utility, denoted u , varies by age, treatment side effects, and disease progression. We draw on the literature measuring health-related quality of life to select calibrated values that affect u . We review the calibrated components and their sources here, and in Section 5 we describe in full how they affect u .

A patient’s utility varies with his age (a , measured in quarters) following [Kind et al. \(1999\)](#) on UK population norms. A patient with local prostate cancer and no treatment side effects receives normalized quarterly utility equal to:

$$Q(a) = 1 - 0.001 \times (a - 80). \quad (8)$$

A patient around age 50 has $Q(200) \approx 0.88$, with utility falling linearly to $Q(\bar{a}) \approx 0.72$. We set the per-quarter discount value δ equal to 0.98.

When a patient receives treatment (A) for his prostate cancer, he may experience side effects for some number of quarters (q) following treatment. We construct a time profile of treatment side effects for ProtecT patients, $L(q)$, using the EPIC scale ([Lane et al., 2022](#)) and scaled to quality adjusted life years (QALYs) with the Patient-Oriented Prostate Utility Scale (PORPUS-U) provided by [Zamora et al. \(2021\)](#). In period q following treatment, a patient’s side effect impact is proportional to $L(q) \in [0, 1]$. At the time of treatment ($q = 0$), $L = 0.039$, and the impact of side effects peaks at $L = 1$ in $q = 4$. The proportional impact of side effects declines to $L = 0.64$ by $q = 24$, and remains approximately constant thereafter. See the appendix for additional information on how we construct $L(q)$. While we assume that all patients experience the same time path of treatment side effects, patients post-treatment utility is additionally scaled by individual preference parameters, described below.

If the patient’s prostate cancer becomes metastatic, we assume that his per-period utility is reduced due to the impact of treatments beyond the prostate (e.g. chemotherapy) and the adverse health impacts of more severe cancer. Specifically, we decrement utility by the factor $\gamma = 0.4$ for individuals with metastatic cancer. This value is consistent with studies by [Kim et al. \(2022\)](#), [Torvinen et al. \(2013\)](#), and [Magnus et al. \(2019\)](#). Once a patient enters

metastasis, we assume that any treatment-related side effects are eliminated, i.e. $L(q) = 0$.

5 Empirical model

We adapt Section 2’s theoretical framework to the details of prostate cancer, its treatment, and the ProtecT trial. Specifically, we make parametric assumptions for model components that we estimate with ProtecT’s publicly available data, and we draw on the calibrated values described in Section 4.2.

To fit the prostate cancer setting, we make a set of additional assumptions on the state vector and its evolution. We add to the state vector an indicator M_a for whether a patient’s prostate cancer is local ($M_a = 0$) or has metastasized ($M_a = 1$). We also track the patient’s age at metastasis, a_M , for patients whose cancer has become metastatic. The patient’s disease state, ω , is a scalar which evolves differently for patients following surveillance (B) versus those who have received treatment (A); we provide additional detail below. Finally, we assume that a patient cannot die from prostate cancer if his disease is local.

5.1 Patient utility

Each patient in our model is diagnosed with prostate cancer, may have received treatment for it, and the cancer could be local or metastatic. A patient’s flow utility depends his current choice, d_a , his state vector χ_a , and his distaste for treatment side effects, α . Flow utility is $u(d_a, M_a, \chi_a, \alpha) \geq 0$, with $u(\cdot) \in (0, 1]$ when the patient is alive and $u(\cdot) = 0$ when the patient is dead.

We construct a live patient’s flow utility from calibrated values analogous to a “QALY,” or quality-adjusted life year. If a patient of age a is alive, has local ($M_a = 0$) cancer, and has not received treatment ($a_A = \emptyset$), he receives the flow utility $u(B, M_a = 0, \chi_a, \alpha) = Q(a)$, where Q is the function described in Section 4.2.

If the patient has received treatment during the current or previous period and his cancer is still local to his prostate, his flow utility is impacted by his disutility from treatment’s side effects. We assume that patients’ side-effect taste values are drawn from $N(\mu_\alpha, \sigma_\alpha^2)$, with μ_α and σ_α as parameters to be estimated. Greater values of α are associated with greater

reductions in utility following treatment. The patient's flow utility is:

$$u(A, M_a = 0, \chi_a, \alpha) = Q(a) [1 - \Lambda(\alpha)L(q)] .$$

In this expression, $\Lambda(\alpha)$ is the logistic CDF, so α 's impact enters u as a $(0, 1)$ scaling factor. $L(q) \in [0, 1]$ is the normalized post-treatment side-effect loss profile, where $q = a - a_A + 1$ denotes quarters since treatment initiation.

If the patient's cancer progresses from local to metastatic ($M_a = 1$), he receives appropriate medical treatment outside of our model. In all periods after the onset of metastatic cancer, the patient's flow utility is scaled-down by $\gamma = 0.4$. We assume that side effects from patients' treatment of localized prostate cancer are not relevant for utility under metastasis, so

$$u(d_a, M_a = 1, \chi_a, \alpha) = \gamma Q(a)$$

In addition to the flow utility components described above, if the patient enrolls in the RCT he receives the payoff shock ϵ_i in the enrollment period, i.e. at age a_0 . We assume that the taste shock from participation is $\epsilon = \exp(\tilde{\epsilon})$, where $\tilde{\epsilon} \sim N(\mu_\epsilon, 1)$. We estimate the parameter μ_ϵ .

5.2 Local disease progression

We assume that the biomarker of local prostate cancer's disease severity, ω_a , tracks PSA scores in integer values from $\omega^{\min} = 3$ to $\omega^{\max} = 30$. In the ProtecT trial, the minimum PSA score is 3, and we assume that PSA growth beyond 30 has no additional impact on disease progression and its associated health risks. In the absence of treatment, ω_a evolves according to $\Pr(\omega_{a+1}|\omega_a)$; see below for our assumptions on ω_a following treatment. Under $d_a = B$, we assume that the disease's severity either remains unchanged between periods or increases by the growth increment Δ^ω up to the cap $\omega^{\max}$. In total, the growth process is:

$$\omega_{a+1} = \begin{cases} \omega_a & \text{with prob. } \exp(-\kappa\omega_a) \\ \min\{\omega_a + \Delta^\omega; \omega^{\max}\} & \text{with prob. } 1 - \exp(-\kappa\omega_a), \end{cases}$$

where Δ_a^ω is drawn from a discrete approximation to the exponential distribution as in Rust (1996, p. 659). We let $J = \omega^{\max} - \omega^{\min}$ and $j \in \{1, 2, \dots, J\}$. The cumulative probability of all possible growth steps (Δ_a^ω) is:

$$C_J(\omega_a) = \int_0^{J+0.5} \lambda_\omega \exp(-\lambda_\omega y) dy,$$

and we use $C_J(\omega_a)$ as a normalizing factor so these probabilities sum to one. We calculate the probabilities of individual growth steps as:

$$\Pr(\Delta_a^\omega) = \begin{cases} \frac{1}{C_J(\omega_a)} \int_0^{0.5} \lambda_\omega \exp(-\lambda_\omega y) dy & \Delta_a^\omega = 0 \\ \frac{1}{C_J(\omega_a)} \int_{j-0.5}^{j+0.5} \lambda_\omega \exp(-\lambda_\omega y) dy & \Delta_a^\omega = 1, \dots, J-1 \\ \frac{1}{C_J(\omega_a)} \int_{J-0.5}^{J+0.5} \lambda_\omega \exp(-\lambda_\omega y) dy & \Delta_a^\omega = J. \end{cases}$$

The underlying continuous exponential density of the increment Δ^ω is characterized by

$$f(\Delta_a^\omega | \omega_a) = \lambda_\omega \exp(-\lambda_\omega \Delta_a^\omega).$$

We specify $\lambda_\omega = \exp(\lambda_{\omega_1} + \lambda_{\omega_2} \omega_a)$, which allows disease progression to advance more quickly as ω_a grows when $\lambda_{\omega_2} < 0$. We estimate the parameters λ_{ω_1} and λ_{ω_2} . Consistent with the cap discussed above, if a realization of Δ_a^ω implies ω_{a+1} would be above $\omega^{\max}$, we set $\omega_{a+1} = \omega^{\max}$.

5.3 Cancer metastasis risk and treatment match value

A patient's ω value has a central role in determining whether his cancer progresses from local to metastatic. If a patient with local cancer is monitoring his disease ($d_a = B$) at age a , his probability of transitioning to metastasis between age a and $a + 1$ is

$$\pi^m(\chi_a, d_a = B) = \Lambda[\lambda_1^m + \lambda_2^m \omega_a]. \quad (9)$$

We estimate the parameters λ_1^m and λ_2^m . If a patient's cancer becomes metastatic, he makes no additional choices in our model.

If the patient receives treatment for his local prostate cancer, the likelihood of metastasis depends on his match value (φ) to treatment. Patients' values of φ are distributed $N(\mu_\varphi, \sigma_\varphi^2)$ in the population, and they are known to the patient prior to the decision to participate in the RCT or undergo any treatment. φ is a source of permanent heterogeneity unobserved by the econometrician. We estimate the parameters μ_φ and σ_φ^2 . A patient with match value φ has a probability $\Lambda(\varphi)$ that his treatment outcome (o) is a success. The uncertainty of treatment success is resolved permanently at the time of treatment.

Consider a patient with the current age a who received treatment at age $a_A < a$, when his PSA score was ω_{a_A} . If this patient's treatment was a success, his probability of transition from local to metastatic cancer between periods a and $a + 1$ is:

$$\pi^m(a, a_A, \omega_{a_A}, d_a = A, o = \text{success}) = \Lambda\left[\lambda_1^m + \lambda_2^m\left(\frac{\omega_{a_A}}{1 + a - a_A}\right)\right].$$

In this function, the probability of metastasis falls rapidly in the periods after treatment.

On the other hand, if treatment is unsuccessful (with probability $[1 - \Lambda(\varphi)]$), the patient's metastasis risk rises weakly over time. We specify that the unsuccessfully-treated patient's risk of metastasis between a and $a + 1$ is a convex combination of two factors. The first factor is his metastasis risk due to his PSA at the time of treatment:

$$p_0(\omega_{a_A}) = \Lambda(\lambda_1^m + \lambda_2^m \omega_{a_A}).$$

p_0 has the same form as the patient's metastasis risk under monitoring, but with ω fixed at its level when the patient received treatment. The second factor is the expected value of his metastasis transition probability at age a if he had remained untreated at age a_A when his PSA score was ω_{a_A} :

$$p_1(a, \omega_{a_A}) = E_\omega\left[\Lambda(\lambda_1^m + \lambda_2^m \omega) | a > a_A, \omega_{a_A}\right].$$

The expectation $E_\omega[\cdot]$ is taken over paths of ω following the stochastic process $\Pr(\omega_{a+1} | \omega_a)$, described above. p_1 captures metastasis risk on a counterfactual path in which the patient had not received treatment at a_A . The total metastasis risk due to failed treatment weighs

the factors p_0 and p_1 by $\rho \in [0, 1]$:

$$\pi^m(a, a_A, \omega_{a_A}, d_a = A, o = \text{failure}) = (1 - \rho) \times p_0(\omega_{a_A}) + \rho \times p_1(a, \omega_{a_A}),$$

where ρ is a parameter to be estimated.

We model the patient's survival from a to $a + 1$ as a two-step process. As argued by [Honoré and Lleras-Muney \(2006\)](#), it is important to allow for competing mortality risks when modeling cancer survival. In the first step, the patient faces non-cancer death risk as a function of age. The probability of dying from a non-prostate cancer source of mortality from a to $a + 1$ is given by:

$$\pi_{a+1}^{nc}(a) = \Lambda\left(\lambda_1^{nc} + \lambda_2^{nc}a\right). \quad (10)$$

In this expression, we estimate the parameters λ_1^{nc} and λ_2^{nc} . In the second step, the patient faces death risk due to prostate cancer. For patients whose cancer remains local, this probability is zero, i.e. $\pi^c(\chi_a) = 0$. For patients with metastatic disease, the probability of a cancer-related death is:

$$\pi^c(\chi_a) = \Lambda\left[\lambda_1^c + \lambda_2^c(a - a_M)\right], \quad (11)$$

where a_M is the age at metastasis. λ_1^{nc} and λ_2^{nc} are parameters to be estimated.

5.4 Value functions

We bring together our modeling framework and the assumptions above to construct the patient's value from treatment, surveillance, and metastasis. We continue with the notation of Section 2 for the value of treatment (A), monitoring (B), and the option to choose between them inside and outside the RCT. We provide some additional detail here to describe how the particular processes and risks of prostate cancer affect these expressions.

If the patient elects to take treatment at age a , he makes no additional choices in the model. The value of treatment is:

$$V_A(\chi_a) = \Lambda(\varphi)V_A(\chi_a, o = \text{success}) + [1 - \Lambda(\varphi)]V_A(\chi_a, o = \text{failure}), \quad (12)$$

where

$$\begin{aligned}
V_A(\chi_a, o) = & u(A, M_a = 0, \chi_a, \alpha) \\
& + \delta[1 - \pi^{nc}(a)] \left\{ \pi^m(\chi_a, d_a = A, o) \times V_M(a + 1, a_M = a + 1) \right. \\
& \left. + [1 - \pi^m(\chi_a, d_a = A, o)] \times V_A(\chi_{a+1}, o) \right\}.
\end{aligned}$$

In this expression for V_A , V_M is the value from entering metastasis:

$$V_M(a, a_M) = u(A, M_a = 1, \chi_a, \alpha) + \delta[1 - \pi^{nc}(a)] \times [1 - \pi^c(\chi_a)] \times V_M(a + 1, a_M). \quad (13)$$

For a patient who chooses to monitor his prostate cancer outside of the RCT, his value function is:

$$\begin{aligned}
V_B(\chi_a, \text{out}) = & u(B, M_a = 0, \chi_a, \alpha) + \delta[1 - \pi_{a+1}^{nc}] \\
& \times \left\{ \pi^m(\omega_a, d_a = B) \times V_M(a + 1, a_M = a + 1) \right. \\
& \left. + [1 - \pi^m(\omega_a, d_a = B)] \times \left(\sum_{\omega_{a+1}} \Pr(\omega_{a+1} | \omega_a) V(\chi_{a+1}, \text{out}) \right) \right\}
\end{aligned}$$

$V(\chi_{a+1}, \text{out})$ is the value of choosing between A and B at age $a + 1$ outside of the RCT and conditional on cancer remaining local ($M_{a+1} = 0$). The value of choosing monitoring in the RCT is similar to its value outside, but with the difference that free choice between B and A occurs only when $\omega_a \geq \underline{\omega}$.

6 Estimation

6.1 Model predictions and the econometric objective function

To estimate our model's parameters, we match empirical moments on disease progression and treatment choices to predictions of the same outcomes from the empirical model. We obtain these predictions using simulation methods.

Conditional on a guess of the model's parameters, θ , we solve our model at every point

in the state space and then simulate outcomes. In particular, we draw a collection of (α, φ) values and then solve backwards the dynamic programming problem of patients in and out of the RCT for every initial $(a_0, \alpha, \omega_0, \varphi)$. We then draw values of ϵ to determine whether patients participate in the RCT. For patients in the RCT and outside of it, we draw sequences of ω values and determine the patients' optimal actions. Where appropriate, we simulate additional outcomes such as treatment success, transitions to metastasis, and mortality. We repeat this process many times to generate frequencies of choices and outcomes, conditional on the guess of θ .

Let $z_k(\theta)$ for $k = \{1, 2, \dots, K\}$ represent a sequence of predicted moments conditional on θ . For each predicted moment, $z_k(\theta)$, we obtain the corresponding empirical frequency, z_k^o from the data sources we describe in Section 4.1. We take each difference between a predicted and empirical moment, $z_k(\theta) - z_k^o$, and stack the differences into the $1 \times K$ vector Z . We estimate the model's parameters by searching over θ values to minimize ZWZ' , where W is a $K \times K$ weighting matrix.

We calculate parameter standard errors using the standard asymptotic variance formula for simulated method of moments, which includes the same weighting matrix W used in point estimation. The covariance combines sampling uncertainty in the published empirical moments with finite-simulation uncertainty estimated from independently redrawn model population. Appendix Section D describes the procedure in detail.

6.2 Parameter estimates

We report parameter estimates and standard errors in Table 1. All the parameters are statistically significant at conventional thresholds, with the exceptions of the σ_φ , which determines the variance of treatment's match value, and $\lambda_{c,2}$ which determines how time since metastasis affects the probability of death from cancer. The model's parameters are difficult to interpret on their own, so we provide some descriptions of what their values imply and how our models' predictions match the data.

In Appendix Table 3 we provide a complete list of the predicted moments (z_k) at the parameter estimates and their observed counterparts (z_k^o), and in Figure 1 we compare the data and outcomes for some key choices and outcomes for patients. In the ProtecT study,

Table 1: Parameter Estimates

Parameter Estimates			
$\lambda_{\omega,1}$	-1.557 (0.284)	$\lambda_{\omega,2}$	0.038 (0.015)
μ_{φ}	2.235 (0.747)	σ_{φ}	0.824 (0.569)
$\lambda_{m,1}$	-7.317 (0.071)	$\lambda_{m,2}$	0.080 (0.011)
μ_{α}	-4.423 (0.068)	σ_{α}	1.595 (0.261)
$\lambda_{nc,1}$	-8.788 (0.297)	$\lambda_{nc,2}$	0.036 (0.004)
$\lambda_{c,1}$	-4.001 (0.333)	$\lambda_{c,2}$	0.032 (0.040)
μ_{ε}	-1.775 (0.242)		
κ	0.021 (0.005)		
ρ	0.974 (0.072)		

62.0% of men who were eligible for randomization accepted the offer to participate in the RCT; at our parameter estimates, our model predicts that 62.3% accept the RCT offer. ProtecT also featured minimal selection on patients’ initial observable characteristics (a_0 and ω_{a_0}), and we obtain a close match on this outcome in our empirical estimates. Our predicted sorting by PSA matches the data well, showing little differences across initial conditions. Our predictions on patient age directionally match key differences across arms, e.g., patients outside the RCT who opt for immediate treatment are younger than those who monitor, but some differences are larger in the predicted values.

In Figure 1 panel (a) we display the empirical and predicted shares of patients initially assigned to B who switched to A , both in the RCT and in the Choice sample. Our model follows the data in generating greater cumulative switching in the RCT than in the Choice data. This feature captures a critical element of patient selection, as the Choice patients are unconstrained by the RCT protocol in their opportunity to switch, yet they do this less frequently. This is driven by the strong selection on distaste for side effects, which we describe further below. See Appendix Figure 8 for a comparison of the ω values at which patients switch from B to A in the RCT and Choice cohorts. While all RCT participants are constrained to switch at $\omega \geq \underline{\omega}$, many Cohort patients switch at lower ω values. Overall, our model matches the empirical switching frequency well for the Choice data. For the first four years after diagnosis, the correspondence between the RCT’s empirical switching share and the model’s prediction is close, but the two measures depart from each other noticeably in the later years of the RCT.


Cumulative patient transitions into metastasis is affected by how soon patients receive treatment. In Appendix Table 3, we report that both RCT and Choice patients who receive A immediately have lower rates of metastasis than patients initially monitoring their prostate cancer. To further discipline the model on important health outcomes, we use reported data from Neal et al. (2020) on metastasis and survival for patients inside and outside the RCT; these moments are not conditioned on an individual’s initial treatment choice. In Figure 1 panel (b) we display overall shares of patients in metastasis in the RCT and Choice samples. Metastasis incidence was generally higher in the Choice data, and this is matched by our model predictions.

The ProtecT data and our estimates show that cancer survival is generally more likely for patients who receive immediate treatment. This is true for both patients randomly assigned to A in the RCT and patients in the Choice sample who select A . In Figure 1 panel (c) we display empirical and model-predicted cancer survival rates, which show that survival is generally higher among RCT participants than those who opt-out.


In Figure 2 panel (a), we display the estimated α population distribution and the conditional α distributions across RCT and Choice cohorts. The conditional distributions of α show substantial selection on side effect distaste across arms. Patients who opt-out of the RCT and take immediate treatment (yellow line) have lower distaste for side effects than the population average (green dashed line). Similarly, patients who opt-out of the RCT and take B (purple line) have substantially greater distaste for treatment’s side effects. By contrast, there is very little systematic separation of cohorts across arms by their φ values, as we display in Figure 2 panel (b). In Appendix Figure 9, we show how the distribution of φ values maps into the distribution of treatment success probabilities, $\Lambda(\varphi)$, with 25th, 50th, and 75th percentile values of 0.85, 0.90, and 0.94, respectively. In Appendix Figure 10 we illustrate the relationship between individuals values of α and φ and their willingness to switch from B to A . Relative to the RCT population, patients in the Choice cohort who initially chose B are more likely to switch to A if they have a low distaste for side effects or a higher expected treatment efficacy.

In Figure 3, we show how patients separate by age, α , and φ values across the treatment and participation options. In this figure, each dot represents an individual patient type for which we simulated an RCT participation decision and subsequent actions for our estimation algorithm. The differently shaded dots indicated patients of different ages, and these show relatively minor differences across cohorts. Differences by φ , measured on the vertical axis, are also minor across cohorts, as would be expected from Figure 2 panel (b). The largest differences among individuals who vary in their α values, measured on the horizontal axis. Patients who opt-out of the RCT and defer immediate treatment (bottom right panel) typically have high values of the distaste from side effects. Those who take treatment right away after declining the RCT (bottom left panel) tend to have low distaste from side effects. In Appendix Figure 11 we provide analogous scatterplots by initial ω_0 , α , and φ values.


Figure 1: Cohort Outcomes Over Time: Data vs. Model



(a) B-to-A switching




(b) Metastasis




(c) Cancer survival

Figure 2: Distributions of α and φ by Cohort




(a) Distribution of α by Cohort




(b) Distribution of φ by Cohort


Figure 3: Simulated Patients by Age, α , and φ , Separated by Cohort




(a) RCT-A



(b) RCT-B



(c) Choice-A



(d) Choice-B

7 What Information Does the RCT Provide?

We use our estimates to describe the impact of endogenous RCT participation on the information generated by the trial. We perform this analysis to characterize selection’s impact given ProtecT’s protocol ($\underline{\omega} = 20$) and for alternative values of $\underline{\omega}$.¹⁰ Our key outcomes are the cancer mortality and metastasis impacts of surveillance (B) versus immediate treatment (A), as recorded through a volunteers’ average treatment effect (VATE) generated by the RCT and population’s average treatment effect (ATE). For an outcome y , the treatment effect is $E[Y_{i,B}(\underline{\omega}) - Y_{i,A}(\underline{\omega})]$ among a relevant population.

We perform our analysis by running a pair of simulations for each $\underline{\omega}$ we consider, including the ProtecT value $\underline{\omega} = 20$. First, to obtain a counterfactual ATE, which is not impacted by voluntary selection, we force all simulated individuals to initiate treatment with A , then initiate treatment with B under $\underline{\omega}$, and then we calculate the differences in outcomes. Different values of $\underline{\omega}$ may yield different values of the ATE, as the protocol impacts the possible treatment actions for patients assigned to B . Second, for a specified value of $\underline{\omega}$ we can allow qualified patients to choose whether to participate in the RCT, and for those who participate we can record the VATE. In this case, the protocol parameter $\underline{\omega}$ affects the treatment options in the experiment (as it does for ATE) and the characteristics of patients who agree to participate in the RCT. Patients’ preferences over treatment side effects affects both willingness to switch from B to A in the experiment and willingness to undergo randomization at all, so selection has impacts on both margins.


In Figure 4, we illustrate the VATE (blue line), ATE (orange line), and difference between the two (black line) for a range $\underline{\omega}$ values. All outcome measures show worse health outcomes for patients who initiate into B than A . For the ProtecT protocol, which has $\underline{\omega} = 20$, we find no difference in VATE and ATE mortality effects 10 years after diagnosis. After 15 years, however, the ATE is larger than the VATE, and this difference is greater at lower values of $\underline{\omega}$. With lower values of $\underline{\omega}$, the RCT attracts participants who would like to switch quickly from B to A , avoiding the negative health consequences of longer waiting for treatment at greater $\underline{\omega}$. A lower value of $\underline{\omega}$ has little impact on patients with a high distaste for treatment

¹⁰Similar calculations could be done for changes to the assignment probability p_A .


side effects. These patients opt-out of an RCT with low $\underline{\omega}$ because they do not want to be randomized to A , and this contributes to a relatively large difference between the ATE and VATE.

Figure 4 also displays differences in metastasis risk. Metastasis is more common than cancer-related death in the ProtecT study, so the outcome measures are larger and the patterns relatively distinct in panels (c) and (d) relative to (a) and (b). The volunteer population has an increased metastasis risk under B than A , but this understates the risk relative to the ATE.


Figure 4: VATE vs. ATE across Alternative Values of $\underline{\omega}$




(a) 10-year cancer mortality



(b) 15-year cancer mortality



(c) 10-year metastasis




(d) 15-year metastasis

We take a closer look at the margins of RCT participation in Figure 5, where we display

a scatterplot of simulated patients by their (α, φ) values. For each (α, φ) pair, we can calculate the value from participating in the RCT versus remaining out of it to freely choose a treatment approach, and we identify the set of patients who would volunteer for the RCT in at least 75% of cases, based on random draws of their participation payoff value ϵ . The solid black line describes the set of patients who meet the 75% participation threshold under ProtecT's $\underline{\omega} = 20$. There is a significant collection of (α, φ) pairs outside of the black boundary with relatively high values of α . These are individuals with a strong enough distaste for treatment's side effects that they generally opt-out of RCT participation. We illustrate the impact of $\underline{\omega}$ on participation by plotting the 75% participation threshold for values of $\underline{\omega} = 16$ (red dashed line) and $\underline{\omega} = 24$ (blue dashed line). With a greater $\underline{\omega}$, the collection of participating patients shrinks. The directions of change are intuitive: patients with lower distaste for side effects or more likely treatment success are less inclined to volunteer when the protocol reduces access to A . With a reduction in $\underline{\omega}$, the margins shift more modestly, with the participation envelope expanding most noticeably at low values of α . These exercises reveal a significant population of high- α individuals who are outside the participation envelopes for all $\underline{\omega}$ values we consider. These patients are the least willing to receive treatment A , and they drive a substantial portion of ATE-VATE differences in the ProtecT population.


Figure 5: Impact of $\underline{\omega}$ on RCT Participation




Notes: Envelope includes types with at least 75% RCT enrollment across preference shock draws.

In Figure 6 we extend the selection analysis of Figure 5 to summarize participation choices for $\underline{\omega}$ values from 4 to 24. As established in Figure 5, Figure 6 shows that patients in the highest tercile of α values are the least likely to participate in the RCT. Side-effect concerns make these patients unlikely to accept randomization. Additionally, Figure 6 shows that these patients' participation choices are relatively unaffected by $\underline{\omega}$; this is due to the patients' reluctance to take treatment A , which implies that less restrictive (i.e., lower $\underline{\omega}$) counterfactual protocols have little additional value. Patients with α in the lowest tercile are also less likely to enroll in the RCT than those with intermediate α values. In contrast to high- α patients, participation by low- α patients is responsive to $\underline{\omega}$ because these individuals are relatively likely to value the opportunity to switch freely from B to A . Thus, the RCT becomes more appealing to them when it is less restrictive.

Figure 6: RCT Participation across ω Values by α and φ Terciles



(a) RCT enrollment share over α tertiles



(b) RCT enrollment share over φ tertiles

8 Conclusions

This paper studies how voluntary participation affects the interpretation of randomized controlled trials (RCTs) comparing treatments that are also available outside the experiment. When patients can obtain their preferred treatment without participating in the RCT, those who agree to randomization tend to be relatively close to indifferent between treatment strategies. As a result, treatment effects estimated among volunteers may differ from those for the full eligible population. An additional complication arises when treatment choices are dynamic and RCT protocols restrict patients' opportunities to switch treatments. These protocols affect both the composition of trial participants and the potential outcomes under the treatments being compared. We develop a framework that jointly models RCT participation, treatment choices outside the experiment, and subsequent treatment decisions. Combining experimental and observational data within this framework allows us to estimate population treatment effects under the observed protocol and evaluate alternative experimental designs.

We apply our framework to prostate cancer treatment using data from the ProtecT study, which includes both randomized participants and eligible patients who declined randomization and selected their own treatment strategies. Our estimates reveal substantial selection on patients' preferences over treatment side effects. Patients who decline randomization and choose immediate treatment tend to have relatively low aversion to side effects, while those

who decline and choose surveillance tend to have high aversion. In contrast, selection on patients' match with treatment is relatively weak, and differences in observable characteristics between volunteers and non-volunteers are modest. These findings illustrate how patients' unobserved preferences can generate substantial sorting into RCT participation even when the experimental sample appears representative along observable dimensions.


Our estimated model allows us to compare treatment effects among RCT volunteers (VATE) with those for the eligible population (ATE). Under the observed ProtecT protocol, the estimated VATE and ATE are approximately equal for 10-year prostate-cancer mortality. However, the ATE exceeds the VATE for 15-year mortality, and the VATE understates the population effect on metastasis at both horizons. We also consider counterfactual protocols that change when patients assigned to surveillance may switch to active treatment. Less restrictive protocols increase participation among patients who value the opportunity to receive active treatment earlier, but have relatively little effect on participation among those with high aversion to treatment side effects. Consequently, changes to switching protocols affect both treatment outcomes and the selection of patients into the RCT. The relationship between VATE and ATE therefore depends on the experimental protocol, the health outcome, and the horizon over which outcomes are measured.

Our findings have broader implications for the interpretation of RCT results and their application outside experimental settings. Even when trial participants resemble the eligible patient population in their observable characteristics, selection on unobserved preferences can generate differences between volunteer and population treatment effects. When treatment choices are dynamic, experimental protocols introduce an additional consideration by affecting both the composition of trial participants and the potential outcomes under the treatments being compared. Data on patients who decline randomization can provide valuable information about the preferences and treatment choices underlying selection into an RCT. Combined with experimental evidence and a model of patient behavior, these data can help recover population treatment effects and evaluate alternative treatment protocols. More broadly, our results illustrate how experimental design can affect both the treatment effects being estimated and the extent of selection into the trial.

Appendix A Side effect impacts

To reflect the impact of treatment-related side effects on patient utility, we construct the function $L(q)$ to provide a normalized loss profile, where $q \geq 0$ denotes quarters since treatment initiation. We construct the side-effect loss profile from mapped Patient-Oriented Prostate Utility Scale (PORPUS-U) trajectories. Let $\widehat{\text{PORPUS}}_{A(q)}$ denote the predicted utility path q quarters after treatment, and let $\widehat{\text{PORPUS}}_{B(q)}$ denote predicted utility path under monitoring in the same quarter had treatment not taken place. We define the treatment-related utility loss as $\tilde{L}(s) = \widehat{\text{PORPUS}}_{B(q)} - \widehat{\text{PORPUS}}_{A(q)}$. PORPUS values may range outside of $[0, 1]$, so we normalize the time path of side effect impacts by its maximum value over q : $L(s) = \frac{L(q)}{\max_q \tilde{L}(q)}$. We display the function $\tilde{L}(q)$ in Figure 7.

Figure 7: Treatment-related loss profile over time, $\tilde{L}$



Appendix B Reweighting and Moment Construction

B.1 Initial Conditions

Table 2 summarizes the distributions used to simulate age and PSA at diagnosis. The bin boundaries are calibrated to the 1,607 men in the ProtecT randomized cohort and the 958 men in the Choice cohort reported in Table 1 of [Neal et al. \(2020\)](#). The RCT assignment probability is treated as a calibrated input rather than an estimation target: each simulated volunteer is assigned to treatment A with the observed probability,

$$\frac{488 + 491}{1,607} = 0.609.$$

Table 2: Distributions Used to Simulate Initial Conditions

Bin	Probability	Randomized cohort	Choice cohort
Panel A: Age at diagnosis (years)			
1	0.25	50.00–58.40	50.00–58.74
2	0.25	58.50–62.40	58.75–62.74
3	0.25	62.50–66.40	62.75–65.74
4	0.25	66.50–69.90	65.75–69.90
Panel B: PSA at diagnosis ($\mu\text{g/L}$)			
1	0.25	[3.00, 3.66]	[3.00, 3.86]
2	0.25	(3.66, 4.63]	(3.86, 4.83]
3	0.18	(4.63, 6.00]	(4.83, 6.20]
4	0.07	(6.00, 6.75]	(6.20, 6.95]
5	0.15	(6.75, 10.00]	(6.95, 10.20]
6	0.10	(10.00, 19.90]	(10.20, 19.90]

Notes: Initial age and PSA are drawn within the selected interval and rounded to the model’s discrete age or PSA grid.

B.2 Construction of the Empirical Moments

The empirical targets are constructed from published ProtecT cohort counts, risk-set tables, incidence rates, and baseline characteristics. We adjust these inputs for attrition, align the published treatment groups with the model, and convert reported rates into probabilities at the model’s evaluation horizons.

Treatment definitions and enrollment. Treatment *A* combines prostatectomy and radiotherapy, while treatment *B* denotes active monitoring. When outcomes are reported separately for prostatectomy and radiotherapy, the treatment-*A* rate is their simple average.

Switching and cohort-level outcomes. Switching, metastasis, and prostate-cancer survival curves are reconstructed from the published numbers at risk and cumulative events. Let n_j denote the number at risk at the start of interval j and C_j the cumulative number of events. We calculate

$$\begin{aligned} d_j &= C_{j+1} - C_j, \\ c_j &= (n_j - n_{j+1}) - d_j, \\ \tilde{n}_j &= n_j - \frac{c_j}{2}, \end{aligned}$$

where d_j is the number of new events, c_j is the inferred number leaving the risk set for other reasons, and $\tilde{n}_j$ adjusts for such attrition within the interval. The cumulative event probability is then constructed by multiplying the interval probabilities of remaining event-free:

$$F(t_{j+1}) = 1 - \prod_{\ell=1}^j \left(1 - \frac{d_\ell}{\tilde{n}_\ell} \right).$$

We apply this procedure to treatment-*B* switching at 2, 4, 6, 8, and 10 years and to the cohort-level metastasis and prostate-cancer survival outcomes reported at three-year intervals.

Arm-specific outcome rates. The published studies report arm-specific all-cause mortality, prostate-cancer mortality, and metastasis rates per 1,000 person-years. After converting

these to annual rates, we define

$$r_{\text{nc}} = r_{\text{all}} - r_{\text{pc}},$$

where r_{nc} is the non-cancer mortality rate. The cumulative probability of death from cause $c \in \{\text{pc}, \text{nc}\}$ over T years is

$$F_c(T) = \frac{r_c}{r_{\text{all}}} (1 - e^{-Tr_{\text{all}}}),$$

and prostate-cancer survival is $1 - F_{\text{pc}}(T)$. Because prostate-cancer death occurs after metastasis in the model, metastasis is constructed with non-cancer death as the competing event:

$$F_{\text{met}}(T) = \frac{r_{\text{met}}}{r_{\text{met}} + r_{\text{nc}}} [1 - e^{-T(r_{\text{met}} + r_{\text{nc}})}].$$

These conversions produce the 10-year arm-specific outcome moments for the RCT and Choice cohorts. The reported five-year RCT prostate-cancer survival proportions enter directly.

Fifteen-year RCT outcomes. To combine the published 10- and 15-year results, we recover the average rate over years 10–15 as

$$r_{10-15} = \frac{15r_{0-15} - 10r_{0-10}}{5} = 3r_{0-15} - 2r_{0-10}.$$

We apply the preceding formulas to this rate over five years and add the resulting increment to the corresponding 10-year cumulative probability. This produces the 15-year metastasis, prostate-cancer mortality, and non-cancer mortality moments.

Appendix C Matched moments and model predictions

Table 3: Predicted and Empirical Moments

Moment No.	Description	Data	Model
<i>Moments regarding RCT enrollment</i>			

Continued on next page

Moment No.	Description	Data	Model
1	Pr(Enrolls in RCT Eligible)	0.620	0.623
<i>Moments from RCT volunteers under trial restrictions</i>			
2	Pr(Switched to A by 24 months after a_0)	0.120	0.118
3	Pr(Switched to A by 48 months after a_0)	0.253	0.283
4	Pr(Switched to A by 72 months after a_0)	0.347	0.458
5	Pr(Switched to A by 96 months after a_0)	0.426	0.602
6	Pr(Switched to A by 120 months after a_0)	0.481	0.702
7	Pr(Metastasis by 120 months $d_{a_0} = B$)	0.056	0.056
8	Pr(Metastasis by 120 months $d_{a_0} = A$)	0.025	0.035
9	Pr(Metastasis by 180 months $d_{a_0} = B$)	0.099	0.078
10	Pr(Metastasis by 180 months $d_{a_0} = A$)	0.051	0.056
11	Pr(Survived cancer at 60 months $d_{a_0} = B$)	0.994	0.995
12	Pr(Survived cancer at 60 months $d_{a_0} = A$)	1.000	0.997
13	Pr(Survived cancer at 120 months $d_{a_0} = B$)	0.983	0.979
14	Pr(Survived cancer at 120 months $d_{a_0} = A$)	0.994	0.988
15	Pr(Survived cancer at 180 months $d_{a_0} = B$)	0.969	0.958
16	Pr(Survived cancer at 180 months $d_{a_0} = A$)	0.974	0.973
17	Pr(Dead by non-cancer cause at 120 months $d_{a_0} = B$)	0.087	0.093
18	Pr(Dead by non-cancer cause at 120 months $d_{a_0} = A$)	0.097	0.093
19	Pr(Dead by non-cancer cause at 180 months $d_{a_0} = B$)	0.198	0.197
20	Pr(Dead by non-cancer cause at 180 months $d_{a_0} = A$)	0.187	0.207
<i>Moments from RCT decliners (Choice cohort)</i>			
21	Pr($d_{a_0} = A$)	0.471	0.490
22	Pr(Switched to A by 24 months after a_0)	0.137	0.127
23	Pr(Switched to A by 48 months after a_0)	0.259	0.261
24	Pr(Switched to A by 72 months after a_0)	0.350	0.377
25	Pr(Switched to A by 96 months after a_0)	0.399	0.467
26	Pr(Switched to A by 120 months after a_0)	0.443	0.526
27	Pr(Metastasis by 120 months $d_{a_0} = B$)	0.056	0.063
28	Pr(Metastasis by 120 months $d_{a_0} = A$)	0.024	0.038
29	Pr(Survived cancer at 120 months $d_{a_0} = B$)	0.979	0.979
30	Pr(Survived cancer at 120 months $d_{a_0} = A$)	0.992	0.986
31	Pr(Dead by non-cancer cause at 120 months $d_{a_0} = B$)	0.092	0.089

Continued on next page

Moment No.	Description	Data	Model
32	Pr(Dead by non-cancer cause at 120 months $d_{a_0} = A$)	0.083	0.076
<i>Supplemental moments on metastasis and cancer survival (Neal et al., 2020)</i>			
33	Pr(Metastasis by 36 months RCT)	0.009	0.012
34	Pr(Metastasis by 72 months RCT)	0.021	0.026
35	Pr(Metastasis by 108 months RCT)	0.034	0.039
36	Pr(Metastasis by 144 months RCT)	0.052	0.052
37	Pr(Metastasis by 36 months CH)	0.014	0.012
38	Pr(Metastasis by 72 months CH)	0.025	0.028
39	Pr(Metastasis by 108 months CH)	0.037	0.045
40	Pr(Metastasis by 144 months CH)	0.053	0.063
41	Pr(Survived cancer at 36 months RCT)	0.999	0.999
42	Pr(Survived cancer at 72 months RCT)	0.996	0.994
43	Pr(Survived cancer at 108 months RCT)	0.993	0.988
44	Pr(Survived cancer at 144 months RCT)	0.985	0.978
45	Pr(Survived cancer at 36 months CH)	0.997	0.999
46	Pr(Survived cancer at 72 months CH)	0.990	0.994
47	Pr(Survived cancer at 108 months CH)	0.985	0.986
<i>Moments on age and PSA initial conditions</i>			
48-50	Initial PSA: RCT, (Q1, Q2, Q3)	(4, 5, 7)	(4, 5, 7)
51-53	Initial PSA: choice $d_{a_0} = B$, (Q1, Q2, Q3)	(4, 5, 7)	(4, 4, 6)
54-56	Initial PSA: choice $d_{a_0} = A$, (Q1, Q2, Q3)	(4, 5, 7)	(4, 6, 8)
57-59	Initial age (in quarters): RCT, (Q1, Q2, Q3)	(234, 250, 266)	(239, 255, 268)
60-62	Initial age (in quarters): choice $d_{a_0} = B$, (Q1, Q2, Q3)	(236, 252, 264)	(236, 253, 267)
63-65	Initial age (in quarters): choice $d_{a_0} = A$, (Q1, Q2, Q3)	(234, 250, 262)	(226, 246, 261)

Appendix D Inference

We calculate standard errors using the standard minimum-distance asymptotic covariance estimator for simulated method of moments (SMM). This section describes the sampling covariance of the empirical moments, the finite-simulation covariance, and the numerical Jacobian used in that calculation. All expressions use the same moment definitions, normalizations, and fixed weighting matrix as the point estimation.

D.1 Covariance of the moment discrepancies

Let $z_k(\theta)$ for $k = \{1, 2, \dots, K\}$ represent a sequence of predicted moments conditional on θ , and let z_k^o represent the corresponding empirical frequency. Let u collect the simulation draws and define the normalized discrepancy for moment k as

$$g_k(\theta, u) = \frac{z_k(\theta, u) - z_k^o}{s_k},$$

where s_k is the empirical moment used to normalize that discrepancy. For the age moments, age is first centered at 200 quarters. Fixed linear contrasts, including the difference between treatment-arm outcomes, and their covariances are transformed using the same normalization. The covariance of the normalized moment vector is

$$\widehat{\Omega} = \widehat{\Omega}_{\text{data}} + \widehat{\Omega}_{\text{sim}}.$$

Both components describe the covariance of the finite-sample moment vector, so the standard sandwich estimator does not apply an additional division by sample size.

Because the empirical targets are published summary statistics rather than statistics calculated from common patient-level data, we reconstruct $\widehat{\Omega}_{\text{data}}$ from the information reported with each target. Enrollment and initial-treatment shares use the binomial variance

$$\widehat{\text{Var}}(\widehat{p}) = \frac{\widehat{p}(1 - \widehat{p})}{n}.$$

For switching, metastasis, and cancer-survival curves, let n_j denote the number at risk and d_j the number of events in interval j . We approximate interval censoring and the effective risk set by

$$c_j = (n_j - n_{j+1}) - d_j, \quad \widetilde{n}_j = n_j - \frac{1}{2}c_j.$$

The resulting life-table estimate and Greenwood covariance are

$$\hat{S}(t_j) = \prod_{\ell \leq j} \left(1 - \frac{d_\ell}{\tilde{n}_\ell}\right),$$

$$\widehat{\text{Cov}}\{\hat{S}(t_i), \hat{S}(t_j)\} = \hat{S}(t_i)\hat{S}(t_j) \sum_{\ell \leq \min(i,j)} \frac{d_\ell}{\tilde{n}_\ell(\tilde{n}_\ell - d_\ell)}.$$

The same covariance applies to cumulative failure, $1 - \hat{S}$, and preserves dependence among dates reported on the same curve.

Published cause-specific event rates are treated as Poisson estimators. If D_c events occur over PY person-years, then

$$\hat{r}_c = \frac{D_c}{PY}, \quad \widehat{\text{Var}}(\hat{r}_c) = \frac{\hat{r}_c^2}{D_c}.$$

The rate is converted to a cumulative probability using the same constant-hazard, competing-risk transformation as the empirical moment,

$$F_c(t) = \frac{r_c}{r_{\text{comp}}} (1 - e^{-tr_{\text{comp}}}),$$

where r_{comp} is the relevant combined hazard. The code averages the surgery and radiotherapy rates used for treatment A and recovers the rate over years 10–15 from

$$r_{10-15} = 3r_{0-15} - 2r_{0-10}.$$

A numerical delta method propagates the Poisson rate variances through these transformations.

The joint covariance of each reported first, second, and third quartile for age and PSA is estimated using 5,000 pseudo-bootstrap samples. Each bootstrap draws the published group sample size from a piecewise-uniform distribution that assigns one quarter of its probability to each interval defined by the reported range and quartiles. Covariance identified within a common life-table curve or nonlinear rate transformation is retained.

To estimate finite-simulation uncertainty, we independently generate $N_{\text{sim}} = 500$ com-

plete model populations at $\hat{\theta}$. Each population contains 1,600 (α, φ) types and 30 clones per type, and all simulated initial conditions and subsequent shocks are redrawn across populations. Using a full population as the replication unit preserves the dependence among clones that share type-level heterogeneity. If $g^{(b)}$ is the normalized moment vector from replication b , then

$$\hat{\Omega}_{\text{sim}} = \frac{1}{N_{\text{sim}} - 1} \sum_{b=1}^{N_{\text{sim}}} (g^{(b)} - \bar{g})(g^{(b)} - \bar{g})'.$$

D.2 Numerical Jacobian and covariance estimator

Let P be the number of estimated parameters and let e_j be the j th unit vector. The j th column of the $K \times P$ Jacobian is estimated using a centered finite difference,

$$\hat{G}_{\cdot j} = \frac{g(\hat{\theta} + h_j e_j, u) - g(\hat{\theta} - h_j e_j, u)}{2h_j}.$$

The primary step is one percent of the parameter's scale. The positive and negative evaluations use common random numbers, which removes simulation noise that is unrelated to the parameter perturbation. The Jacobian and final covariance are expressed in the reported parameter units. When an unconstrained numerical coordinate is used for a constrained parameter, the chain rule maps the covariance back to the reported scale.

Let $W = \text{diag}(w_1, \dots, w_K)$, where w_k are fixed, block-normalized, moment-specific weights used in the estimation criterion. In practice, these moments weights are prespecified multipliers that provide additional emphasis on economically important moments.¹¹ Define

$$\begin{aligned}\hat{A} &= \hat{G}' W \hat{G}, \\ \hat{B} &= \hat{G}' W \hat{\Omega} W \hat{G}.\end{aligned}$$

¹¹For moment number k , as ordered in Table 3, the baseline weights are $(6, \frac{5}{11}, \frac{1}{11}, \frac{12}{11}, \frac{1}{11}, \frac{20}{11}, \frac{2}{7}, \frac{2}{7}, \frac{9}{7}, \frac{3}{7}, \frac{1}{15}, \frac{1}{15}, \frac{1}{15}, \frac{1}{15}, 1, \frac{4}{3}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{24}{11}, \frac{14}{11}, \frac{1}{11}, \frac{15}{11}, \frac{1}{11}, \frac{18}{11}, \frac{5}{7}, \frac{4}{7}, 1, \frac{1}{3}, \frac{1}{6}, \frac{1}{6}, \frac{1}{14}, \frac{1}{14}, \frac{1}{14}, \frac{1}{14}, \frac{1}{14}, \frac{1}{14}, \frac{1}{14}, \frac{1}{14}, \frac{1}{15}, \frac{1}{15}, \frac{1}{15}, \frac{1}{15}, \frac{1}{15}, \frac{1}{15}, \frac{1}{15}, \frac{1}{15}, \frac{1}{15}, \frac{5}{9}, \frac{1}{18}, \frac{1}{18}, \frac{7}{9}, \frac{1}{18}, \frac{1}{18}, \frac{8}{9}, \frac{8}{9}, \frac{5}{9}, \frac{5}{9}, \frac{5}{9}, \frac{5}{9}, \frac{5}{9}, \frac{5}{9}, \frac{10}{9}, \frac{10}{9}, \frac{10}{9}, 2)$.

The covariance matrix and parameter standard errors are


$$\widehat{V}_\theta = \widehat{A}^{-1} \widehat{B} \widehat{A}^{-1},$$

$$\widehat{\text{SE}}(\widehat{\theta}_j) = \sqrt{\{\widehat{V}_\theta\}_{jj}}.$$


The reported inference therefore incorporates both uncertainty in the published targets and finite model simulation. It conditions on calibrated inputs that are not estimated and does not incorporate uncertainty from model misspecification.

Appendix E Estimation results

Figure 8: Distribution of ω at Switch to A by Cohort



(a) RCT Cohort



(b) Choice Cohort

Figure 9: Patient Success Probabilities







Figure 10: Patient Characteristics and Differences in Actions under B




(a) RCT-B: α distribution



(b) RCT-B: φ distribution




(c) Choice-B: α distribution




(d) Choice-B: φ distribution


Figure 11: Simulated Patients by ω_0 , α and φ , Separated by Cohort




(a) RCT-A



(b) RCT-B



(c) Choice-A



(d) Choice-B

References

- Andrews, Isaiah and Emily Oster**, “A Simple Approximation for Evaluating External Validity Bias,” *Economics Letters*, 2019, 178, 58–62. [7](#)
- Athey, Susan, Raj Chetty, and Guido Imbens**, “Combining Experimental and Observational Data to Estimate Treatment Effects on Long Term Outcomes,” 2020. [1](#)
- Basu, Anirban, James J Heckman, and Sergio Urzua Salvador Navarro-Lozano**, “Use of instrumental variables in the presence of heterogeneity and self-selection: an application to treatments of breast cancer patients,” *Health Economics*, November 2007, 16 (11), 1133–1157. [1](#)
- Bill-Axelson, Anna, Lars Holmberg, Hans Garmo, Kimmo Taari, Christer Busch, Stig Nordling, Michael Häggman, Swen-Olof Andersson, Ove Andrén, Gunnar Steineck, Hans-Olov Adami, and Jan-Erik Johansson**, “Radical Prostatectomy or Watchful Waiting in Prostate Cancer - 29-Year Follow-up,” *New England Journal of Medicine*, 2018, 379, 2319–2329. [3](#), [4.1](#)
- Carroll, P.H. and J.L. Mohler**, “NCCN guidelines updates: prostate cancer and prostate cancer early detection,” *J Natl Compr Canc*, 2018, 16 (5S), 620–623. [3](#)
- Chan, Tat Y. and Barton H. Hamilton**, “Learning, Private Information and the Economic Evaluation of Randomized Experiments,” *Journal of Political Economy*, December 2006, 114 (6), 997–1040. [4](#)
- Chen, RC, JA Clark, and JA Talcott**, “Individualizing quality-of-life outcomes reporting: how localized prostate cancer treatments affect patients with different levels of baseline urinary, bowel, and sexual function,” *Journal of Clinical Oncology*, 2009, 27 (24), 3916–3922. [3](#)
- Epstein, MM, G Edgren, JR Rider, LA Mucci, and HO Adami**, “Temporal trends in cause of death among Swedish and US men with prostate cancer,” *Journal of the National Cancer Institute*, 2012, 104 (17), 1335–1342. [3](#)

Farah, Wigdan H, Mouaz Alsawas, Maria Mainou, Fares Alahdab, Magdoleen H Farah, Ahmed T Ahmed, Essa A Mohamed, Jehad Almasri, Michael R Gionfriddo, Ana Castaneda-Guarderas, Khaled Mohammed, Zhen Wang, Noor Asi, Craig N Sawchuk, Mark D Williams, Larry J Prokop, M Hassan Murad, and Annie LeBlanc, “Non-pharmacological treatment of depression: a systematic review and evidence map,” *BMJ Evidence Based Medicine*, December 2016, *21* (6), 214–221. [1](#)

Hamdy, Freddie C., Jenny L. Donovan, J. Athene Lane, Chris Metcalfe, Michael Davis, Emma L. Turner, G.J. Martin Richard M. Young, Eleanor Walsh, R.J. Bryant, Prasad Bollina, Andrew Doble, Alan Doherty, David Gillatt, O. Gnanapragasam, O. Hughes, Roger Kockelbergh, Howard Kynaston, Alan Paul, E. Paez, Philip Powell, Derek J. Rosario, Edward Rowe, Malcolm Mason, James Catto, Tim J. Peters, Jon Oxley, N.J. Williams, John Staffurth, and David E. Neal, “Fifteen-Year Outcomes after Monitoring, Surgery, or Radiotherapy for prostate Cancer,” *New England Journal of Medicine*, 2023, *388* (17), 1547–1558. [3](#), [4.1](#)

—, —, —, —, Malcolm Mason, Chris Metcalfe, Peter Holding, Michael Davis, Tim J. Peters, Emma L. Turner, Richard M. Martin, Jon Oxley, Mary Robinson, John Staffurth, Eleanor Walsh, Prasad Bollina, James Catto, Andrew Doble, Alan Doherty, David Gillatt, Roger Kockelbergh, Howard Kynaston, Alan Paul, Philip Powell, Stephen Prescott, Derek J. Rosario, Edward Rowe, and David E. Neal, “10-Year Outcomes after Monitoring, Surgery, or Radiotherapy for Localized Prostate Cancer,” *New England Journal of Medicine*, 2016, *375* (15), 1415–1424. PMID: 27626136. [1](#), [3](#), [4.1](#)

Heckman, James J., “Randomization and Social Policy Evaluation,” in Charles F. Manski and Irwin Garfinkel, eds., *Evaluating Welfare and Training Programs*, Cambridge, MA: Harvard University Press, 1992, chapter 5, pp. 201–230. [1](#)

— and Edward Vytlačil, “Local instrumental variables and latent variable models for identifying and bounding treatment effects,” *Proceedings of the National Academy of Sciences*, 2001, *96* (8), 4730–4734. [1](#)

- **and** — , “Structural Equations, Treatment Effects, and Econometric Policy Evaluation,” *Econometrica*, May 2005, *73* (3), 669–738. [1](#), [7](#)
- Honoré, Bo E and Adriana Lleras-Muney**, “Bounds in Competing Risks Models and the War on Cancer,” *Econometrica*, 2006, *74* (6), 1675–1698. [5.3](#)
- Huhn, Maximilian, Adriani Nikolakopoulou, Johannes Schneider-Thoma, Marc Krause, Myrto Samara, Natalie Peter, Thomas Arndt, Lio Bäckers, Philipp Rothe, Andrea Cipriani, John Davis, Georgia Salanti, and Stefan Leucht**, “Comparative efficacy and tolerability of 32 oral antipsychotics for the acute treatment of adults with multi-episode schizophrenia: a systematic review and network meta-analysis,” *Lancet*, 2019, *394*, 939–951. [1](#)
- Jackson, Jeffrey L, Elizabeth Cogbill, Rafael Santana-Davila, Christina Eldredge, William Collier, Andrew Gradall, Neha Sehgal, and Jessica Kuester**, “A Comparative Effectiveness Meta-Analysis of Drugs for the Prophylaxis of Migraine Headache,” *PLoS ONE*, 2015, *10* (7), e0130733. [1](#)
- Johnson, RL, SL Kopp, CM Burkle, CM Duncan, AK Jacob, PJ Erwin, MH Murad, and CB Mantilla**, “Neuraxial vs general anaesthesia for total hip and total knee arthroplasty: a systematic review of comparative-effectiveness research,” *British Journal of Anaesthesia*, 2016, *116* (2), 163–176. [1](#)
- Katterman, Shawn N, Brighid M Kleinman, Megan M Hood, Lisa M Nackers, and Joyce A Corsica**, “Mindfulness meditation as an intervention for binge eating, emotional eating, and weight loss: A systematic review,” *Eating Behaviors*, 2014, *15*, 197–204. [1](#)
- Kelly, John F, Keith Humphreys, and Marica Ferri**, “Alcoholics Anonymous and other 12-step programs for alcohol use disorder,” *Cochrane Database of Systematic Reviews*, 2020, *3*. [1](#)

- Kim, Seon-Ha, Minsu Ock, Min-Woo Jo, and Sungchan Park**, “Estimation of Utility Weights for Prostate-Related Health States in Korea,” *Journal of Preventive Medicine and Public Health*, 2022, *55* (3), 243–252. [4.2](#)
- Kind, Paul, Geoffrey Hardman, and Susan Macran**, “UK Population Norms for EQ-5D,” *Centre for Health Economics, University of York*, 1999. [4.2](#)
- Lane, Janet Athene, Jenny L. Donovan, Grace J. Young, Michael Davis, Eleanor I. Walsh, Kerry N. L. Avery, Jane M. Blazeby, Malcolm D. Mason, Richard M. Martin, Tim J. Peters, Emma L. Turner, Julia Wade, Prasad Bollina, James W. F. Catto, Alan Doherty, David Gillatt, Vincent Gnanapragasam, Owen Hughes, Roger Kockelbergh, Howard Kynaston, Jon Oxley, Alan Paul, Edgar Paez, Derek J. Rosario, Edward Rowe, John Staffurth, David E. Neal, Freddie C. Hamdy, and Chris Metcalfe**, “Functional and Quality of Life Outcomes of Localised Prostate Cancer Treatments (Prostate Testing for Cancer and Treatment [ProtecT] Study),” *BJU International*, 2022, *130* (3), 370–380. [4.2](#)
- List, John A**, “Two Selection Problems, One Bias Term: Experimental Sign-up Is Treatment Choice in Disguise,” *National Bureau of Economic Research*, 2026. [1](#), [2.4](#), [7](#)
- Litwin, MS, RD Hays, A Fink, PA Ganz, B Leake, GE Leach, and RH Brook**, “Quality-of-life outcomes in men treated for localized prostate cancer,” *Journal of the American Medical Association*, 1995, *273* (2), 129–135. [3](#)
- Loeb, S, N Byrne, DV Makarov, H Lepor, and D. Walter**, “Use of conservative management for low-risk prostate cancer in the Veterans Affairs integrated health care system from 2005-2015,” *JAMA*, 2018, *319* (21), 2231–2233. [3](#)
- Magnus, Anne, Wanrudee Isaranuwatthai, Cathrine Mihalopoulos, Victoria Brown, and Rob Carter**, “A Systematic Review and Meta-Analysis of Prostate Cancer Utility Values of Patients and Partners Between 2007 and 2016,” *MDM Policy & Practice*, 2019, *4* (1), 2381468319852332. [4.2](#)

- Mahal, Brandon A, Santino Butler, Idalid Franco, Daniel E Spratt, Timothy R Rebbeck, Anthony V D’Amico, and Paul L Nguyen**, “Use of Active Surveillance or Watchful Waiting for Low-Risk Prostate Cancer and Management Trends Across Risk Groups in the United States, 2010-2015,” *JAMA*, 2019, *321* (7), 704–706. [3](#)
- Malani, Anup**, “Patient enrollment in medical trials: Selection bias in a randomized experiment,” *Journal of Econometrics*, 2008. [1](#)
- Neal, David E, Chris Metcalfe, Jenny L Donovan, J Athene Lane, Michael Davis, Grace J Young, Susan J Dutton, Eleanor I Walsh, Richard M Martin, Tim J Peters, Emma L Turner, Malcolm Mason, Richard Bryant, Prasad Bollina, James Catto, Alan Doherty, David Gillatt, Vincent Gnanapragasam, Peter Holding, Owen Hughes, Roger Kockelbergh, Howard Kynastonm, Jon Oxley, Alan Paul, Edgar Paez, Derek J Rosario, Edward Rowe, John Staffurth, Doug G Altman, Freddie C Hamdya, and for the ProtecT Study Group**, “Ten-year Mortality, Disease Progression, and Treatment-related Side Effects in Men with Localised Prostate Cancer from the ProtecT Randomised Controlled Trial According to Treatment Received,” *European Urology*, November 2020, *77*, 320–330. [3](#), [4.1](#), [6.2](#), [B.1](#)
- NHS England Digital**, “Cancer Registration Statistics, England, 2023.” [3](#)
- Rust, John**, “Numerical Dynamic Programming in Economics,” in H.M. Amman, D.A. Kendrick, and J. Rust, eds., *Handbook of Computational Economics*, Vol. 1, Amsterdam: Elsevier Science B. V., 1996, chapter 14, pp. 619–729. [5.2](#)
- Sianesi, Barbara**, “Evidence of randomisation bias in a large-scale social experiment: The case of ERA,” *Journal of Econometrics*, May 2017, *198* (1), 41–64. [1](#)
- Siegel, Rebecca L., Tyler B. Kratzer, Nikita Sandeep Wagle, Hyuna Sung, and Ahmedin Jemal**, “Cancer Statistics, 2026,” *CA: A Cancer Journal for Clinicians*, 2026, *e70043*, 1–34. [3](#)
- Thomas, J Graham, Hollie A Raynor, Dale S Bond, Anna K Luke, Chelsi C Cardoso, Gary D Foster, and Rena R Wing**, “Weight Loss in Weight Watchers On-

line with and without an Activity Tracking Device Compared to Control: A Randomized Trial,” *Obesity*, 2017, *25*, 1014–1021. [1](#)

Torvinen, Saku, Niilo Färkkilä, Harri Sintonen, Tiina Saarto, Risto P. Roine, and Kimmo Taari, “Health-Related Quality of Life in Prostate Cancer,” *Acta Oncologica*, 2013, *52* (6), 1094–1101. [4.2](#)

Wilt, TJ, MK Brawer, KM Jones, MJ Barry, WJ Aronson, S Fox, JR Gingrich, JT Wei, P Gilhooly, PM Grob, I Nsouli, P Iyer, R Cartagena, G Snider, C Roehrborn, R Sharifi, W Blank, P Pandya, GL Andriole, D Culkin, T Wheeler, and Prostate Cancer Intervention versus Observation Trial (PIVOT) Study Group, “Radical prostatectomy versus observation for localized prostate cancer,” *New England Journal of Medicine*, 2012, *367* (3), 203–213. [3](#), [4.1](#)

Yebyo, Henock G, Hélène E Aschmann, Marco Kaufmann, and Milo A Puhan, “Comparative effectiveness and safety of statins as a class and of specific statins for primary prevention of cardiovascular disease: A systematic review, meta-analysis, and network meta-analysis of randomized trials with 94,283 participants,” *American Heart Journal*, 2019, *210*, 18–28. [1](#)

Yunusa, Ismaeel, Adnan Alsumali, Asabe E Garba, Quentin R Regestein, and Tewodros Egualé, “Assessment of Reported Comparative Effectiveness and Safety of Atypical Antipsychotics in the Treatment of Behavioral and Psychological Symptoms of Dementia: A Network Meta-analysis,” *JAMA Network Open*, 2019, *2* (3), e190828. [1](#)

Zamora, Víctor, Olatz Garin, Yolanda Pardo, Àngels Pont, Cristina Gutiérrez, Patricia Cabrera, Francisco Gómez-Veiga, José Ignacio Pijoan, Mark S. Litwin, and Montse Ferrer, “Mapping the Patient-Oriented Prostate Utility Scale From the Expanded Prostate Cancer Index Composite and the Short-Form Health Surveys,” *Value in Health*, 2021, *24* (11), 1676–1685. [4.2](#)