Correlated Drug Action as a Baseline Additivity Model for Combination Cancer Therapy in Patient Cohorts and Cell Cultures

Adith S. Arun[1,2], **Sung-Cheol Kim**[1,3], **Mehmet Eren Ahsen**[4], **Gustavo Stolovitzky**[1,†]

[1] IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, USA
[2] Johns Hopkins University, Baltimore, Maryland 21218, USA
[3] Psychogenics, Inc, Paramus, New Jersey, 07652, USA
[4] Gies College of Business, University of Illinois at Urbana-Champaign, Illinois 61820, USA

[†] To whom correspondence should be addressed: gustavo@us.ibm.com

Abstract

Identifying and characterizing the effect of combination cancer therapies is of paramount importance in cancer research. The benefit of a combination can either be due to inherent heterogeneity in patient populations or because of molecular synergy between the compounds given in combination, usually studied in cell culture, or both. To shed light and help characterized combinations and their enhanced benefits over single therapies, we introduce Correlated Drug Action (CDA) as a baseline additivity model. We formulate the CDA model as a closed-form expression, which lends itself to be scalable and interpretable, both in the temporal domain (tCDA) to explain survival curves, and in the dose domain (dCDA), to explain dose-response curves. CDA can be used in clinical trials and cell culture experiments. At the level of clinical trials, we demonstrate tCDA's utility in explaining the benefit of clinical combinations, identifying non-additive combinations, and cases where biomarkers may be able to decouple the combination into monotherapies. At the level of cells in culture, dCDA naturally embodies null models such as Bliss additivity and the Highest Single Agent model as special cases, and can be extended to be sham combination compliant. We demonstrate the applicability of dCDA in assessing non-additive combinations and doses. Additionally, we introduce a new synergy metric, Excess over CDA (EOCDA), that incorporates elements of Bliss additivity and dose equivalence concepts in the same measure. CDA is a novel general framework for additivity at the cell line and patient population levels and provides a method to characterize and quantify the action of drug combinations.

Keywords: Combination Cancer Therapy, Additivity Models

Introduction

Combination therapy, the principle of treating a patient with multiple treatments either simultaneously or in succession, is a widely used strategy for treating conditions ranging from HIV/AIDS to cancer [1–

4]. In the 1960's Frei and Freirich reasoned that cancer cells could mutate and become resistant to one chemotherapy but were less likely to become resistant to multiple different anti-cancer chemotherapy agents [5]. Since that pioneering work, combination therapies have become the de facto approach for treating cancer. Thus, it is of paramount importance to quantitatively characterize the effects of drugs acting in combination. Additivity models, that is, models that assume the combination effect results from the addition of the effects of the monotherapies, can serve as baselines for the expected effects of drugs applied in combination [6], for example as reference to determine if a combination is synergistic.

There are fundamentally two different levels at which the effect of combinations have been studied: cultured cells or patient populations. Cell lines and organoids fit into the former category whereas Patient Derived Xenografts (PDXs) mice models and human clinical trials are examples of the latter. We will refer to additivity models at the level of cultured cells as dose-space models since the primary variable to change is the dose of the monotherapies in the combination. Similarly, models describing patient populations will be referred to as temporal models given that the main endpoint is survival times.

Recently, a framework based on the principle of Independent Drug Action (IDA) has been proposed to describe additive effects of drug combinations in survival analysis in patient cohorts [7]. IDA is a temporal model that asserts independence between the individual monotherapies that constitute the combination and posits that the survival of the patient treated with a combination of drugs is equal to the survival that the patient would have had, if they had been treated with the monotherapy with larger survival time for that patient. This is in essence an independence-based additivity model that can be useful to model combinations in patients under the assumption that the mode of action of one of the drugs does not affect the mode of action of the other drugs. Interestingly, the IDA model brings patient to patient heterogeneity to the forefront of additivity models in the temporal space. In this paper we extend upon the previous work by Sorger and Palmer [7]. As one of the essential features in Sorger and Palmer [7] is the consideration of the potentially correlated action of two drugs on the same patient, we will call our extension of IDA the temporal Correlated Drug Action (tCDA) additivity model. This is to emphasize that even though the monotherapies in a combination may not causally affect each other, they may be correlated through confounding factors. The nature of this correlation is difficult to ascertain in general, and many models can be built that have a given correlation. Our model makes some assumptions about the way in which the correlation operated that allow us to formulate a closed-form and interpretable expression for the effect of the combinations in terms of the effects of the monotherapies (See Supplement). We also build out other correlation models based on simulations and show their relation to the closed-form model. We developed methodology to estimate the probability that the tCDA model can be an explanation for the effect of the combination.

Translating these temporal ideas around dependence to the dose-space, we propose the dCDA model that describes the effect of combinations in cell cultures as a function of the dosages of each monotherapy. Additivity models at the level of cell cultures are much more developed [8–10] than those describing survival times of combinations in patient populations [11]. Amongst the most-used reference dose-space additivity models are Bliss additivity [12], Loewe additivity [13], the Dose Equivalence Principle

[14], and Highest Single Agent (HSA) [15]. We will present two formulations of our dCDA additivity model, one that naturally embeds Bliss and HSA models at two extremes of correlations, and another which incorporates the principle of sham combination compliance to the dCDA, which we will call sham compliant dCDA.

CDA, both temporal and dose-space, are based on the assumption of no direct effect of one drug on the mechanism of the other. Broadly, non-interaction means that the two monotherapies do not affect one another at the cellular, tissue, organ, or organism level. For the temporal model, tCDA allows for there to be a correlation between the survival times under either monotherapy for each patient. This correlation cannot be directly measured, and can only be inferred in terms of the effects it produces in a population. Likewise, for the dose-space model, the dCDA accounts for a correlation between the lethal doses under either monotherapy for each cell in a cultured population of cells. The range of correlation values ($[-1, 1]$) allows for a gamut of possible outcomes for the effect of a combination under dCDA. For the dose-space CDA, at different values of correlation, the model describes a model closer to Bliss additivity (when correlation is close to 0) or a HSA model (when correlation is closer to 1). We must make clear that the correlation described in either model cannot be measured experimentally because it is not possible to treat a given patient or cell with one monotherapy and then reset to the pre-initial monotherapy conditions and administer the other monotherapy. Thus, the correlation is a latent variable that must be estimated from the data and can inform us of the underlying biology of the combination (e.g. whether a biomarker can stratify the patient population into those that respond better to one monotherapy or the other).

We need not constrain ourselves to two drug combinations. The CDA framework allows for a combination to be composed of three drugs, say A, B, and C, so long as there is data available for drug C and for drug A and B together. Naturally, this logic can be extended to a combination composed of any number of drugs. Thus, the term monotherapy can be applied liberally and need not only represent a single drug. In many of the clinical trials analyzed in this work, there are three or four drugs that form the combination as such an approach has been historically successful [16–19].

We apply the temporal CDA null additivity model to oncology clinical trial data in order to demonstrate its ability to explain the effect of clinical combination therapies and test for non-additive interactions at work within these combinations. Furthermore, we estimate, as a parameter in the model, the correlation of the combination which has potential biological implications. Notably, the correlation informs whether there may exist a biomarker that can decouple the combination into its constituent monotherapies while realizing the benefit of the combination.

Rather than introducing an entirely new model in an area that has seen much debate [20], the dose-space CDA null additivity model seamlessly marries different individual concepts (Bliss, HSA, the dose equivalence principle and sham combination compliance) and creates a spectrum of non-interaction possibilities between these established models. This semi-unification occurs rather naturally from the basic premises of the approach used to develop the CDA model. We show that the dose-space CDA model

is useful in assessing whether there are non-additive effects in a given combination tested on cultured cells. Additionally, we introduce a new performance metric Excess over CDA (EOCDA), and validate the applicability of the commonly used Excess over Bliss (EOB) metric in the tested combinations.

Results

Temporal Correlated Drug Action as a model for understanding combination therapy in human clinical trials

As a representative example of the tCDA model, we can examine the combination of chemotherapy and Trastuzumab (Herceptin). Chemotherapy consisted of anthracycline plus cyclophosphamide for patients who had never received anthracycline before or paclitaxedl for patients who had received adjuvant anthracycline [21]. This therapy is a current treatment for HER-2 overexpressing breast cancers. The Herceptin and chemotherapy combination shows a clinical benefit over either individual monotherapy with respect to progression-free survival (PFS) time (Fig. 1A). At around 12 months, the benefit of the combination reduces to that of the Herceptin treated group. The tCDA model predicts the PFS curve $S_{AB}(t)$ of the combination of the two drugs A and B (in this case chemotherapy and Trastuzumab) in terms of the PFS curves of each of the monotherapies $S_A(t)$ and $S_B(t)$ in the form prescribed by the mathematical expression

$$S_{AB}(t,\rho) = \\ = \begin{cases} S_{\mathsf{max}}(t) + S_{\mathsf{min}}(t)(1 - S_{\mathsf{max}}(t))(1-\rho) & \text{if } 0 \leq \rho \leq 1 \\ (S_A(t) + S_B(t) - S_A(t)S_B(t))(1-|\rho|) + \min(1, S_A(t) + S_B(t))|\rho| & \text{if } -1 \leq \rho_s \leq 0 \end{cases} \quad (1)$$

where $S_{\mathsf{max}}(t) = \max(S_A(t), S_B(t))$, $S_{\mathsf{min}}(t) = \min(S_A(t), S_B(t))$, and ρ is the correlation over all patients of the times that each patient would have survived if treated with drug A and B independently. (See Supplement for the derivation of Eq.(1).) As we vary the parameter ρ in Eq.(1) using as $S_A(t)$ and $SB(t)$ the observed PFS curves of the chemotherapy and Trastuzumab respectively, we span the range of the tCDA predictions forming the cone of possibilities shown in Fig. 1B, where each color corresponds to one value of correlation.

The tCDA model embodied in Eq. (1) is based on the premise that there is a strong correlation between the survival times for each drug for a fraction ρ of the patients, and for the other fraction (1-ρ) their survival times are uncorrelated (See Supplement). Fig. 1C shows different simulations (using the Coin method described in the Supplement) of survival times for different overall correlation in the patient population. Each point in this Figure is a patient, and the y-axis and x-axis values are the survival times such patient would have had if treated with Trastuzumab or chemotherapy respectively. THe distribution of x-axis and y-axis values in all subfigures follow $S_A(t)$ and $S_B(t)$ curves respectively, and the correlations alluded to in the figure as Spearman's rank correlations. A value of Spearman's

correlation equal to 1 traces the Dose Equivalence Curve between A and B (see Supplement) while a value of -1 associates the survival times of patients to drugs A and B in reverse order. A spearman correlation of zero assigns survival times to A and B independently. Intermediate Spearman's correlation coefficients will have fewer points on the Spearman's correlation 1 or -1 curve - a move away from order and towards disorder. From this joint distribution of patient survival times and taking the maximum value at each ordered pair, we can construct the empirical estimate of the PFS curve under temporal CDA at a given Spearman's correlation. This simulation process is mathematically equivalent to the closed-form expression given in Eq.(1) (See Supplement and Fig. S6) which we will use throughout the paper. We will estimate the tCDA predicted PFS curve of the combination by varying the parameter ρ between $[-1, 1]$ and choosing the value that minimizes the error between the estimated and observed combination PFS curves.

Returning to the Herceptin and chemotherapy trial, the optimal estimate for the Spearman's correlation is 0.03 and its corresponding 95% confidence interval is $[-0.07, 0.13]$. Interestingly, a near-zero ρ, as we have here, collapses the temporal CDA model into a form analogous to that of Bliss independence. More importantly, the relative value of ρ has a potential biological meaning: a near-zero or zero ρ represents a combination that maximizes the effect of its constituent monotherapies with respect to the underlying heterogeneity in the treated patient population (See Methods) even if there is no direct interaction between the drugs. Practically, this suggests that it may not be useful to look for stratify metastatic HER-2 overexpressing breast cancer patients into those that respond better to Trastuzumab and those that respond better to chemotherapy because the heterogeneity in the patients' tumors (e.g. tumor microenvironments, metastatic lineages, clonal populations) render it a near-futile task. It may be better to administer the combination to all such patients in comparison to trying to assign each patient to the monotherapy that will work best for them. In this scenario, we are fundamentally unable to decide for a given patient whether Trastuzumab or chemotherapy is better for them and in order to reduce the guesswork in the process we are forced to give the combination for each patient. Contrarily, when the absolute value of the Spearman's correlation approaches one (Fig. 2A), it may be possible to stratify the patient population such that we need not administer the combination and only assign one of the constituent monotherapies to each individual patient based on a biomarker panel. Examples of such include 5-FU and Oxaliplatin in advanced pancreatic cancer, Interferon Alfa and Temsirolimus in advanced renal cell carcinoma, Irinotecan Bevacizumab and Panitunumab in advanced colorectal cancer, and other combination chemotherapy regimens (Supp. Figs. 3 - 5). Of course, we must be careful to not assign meaning or further analyze the results without confirming that the temporal CDA model sufficiently describes the effect of the combination.

We can observe that the resulting estimate for the combination PFS curve under temporal CDA follows the true combination rather well (Fig. 1D). Empirically, using a modified Kolmogorov-Smirnov test statistic (See Supplement), we fail to reject the null hypothesis that the tCDA model sufficiently describes the observed combination (p-value of 0.46). The tCDA model can explain the benefit of this combination from its respective monotherapy survival curves. In this sense it can be hypothesized that

this combination works according to tCDA additivity.

Temporal Correlated Drug Action represents a null additivity model for discerning non-additive combinations

For 18 different combinations tested in clinical trials, we estimate the optimal tCDA predicted survival curve and assess whether tCDA is the appropriate framework to apply (Fig. 2A). Using the tCDA framework as a null model allows us to categorize clinically tested combinations as either tCDA or non-CDA compatible via a formal hypothesis test. The null hypothesis is that temporal CDA adequately describes the observed combination. We use a p-value threshold of 0.01, and this threshold reflects our motivation that we should fail to reject the null (additivity) only when the combination is truly non-additive. In 61% of the tested combinations, the temporal CDA model adequately describes the effect of the combination on a patient population as noted by a p-value greater than 0.01 (Fig. 2A). The estimated combination PFS curve closely tracks the observed combination in such cases (Figs. 1D, 2B-F; Supp. Figs. 3C, 3D, 4A, 5A, 5C). Notably, the tCDA model can handle combinations composed of more than two drugs so long as there is some partition of the drugs within the combination into two distinct monotherapy PFS curves (e.g., Fig. 2F).

Observed combinations that closely follow the best-performing monotherapy represent situations where it is most beneficial to treat patients only with the consistently better performing monotherapy since there are fewer concerns around drug-drug interactions, toxicity, and side effects that accompany combination therapies. An example of this is the combination of Interferon Alfa and Temsirolimus in advanced renal cell carcinoma (Supp. Fig. 3C) because Temsirolimus is consistently better performing monotherapy across time and should replace the combination. In Fig. 2E immediately after 15 months, the tCDA PFS estimate follows the Paclitaxel and Carboplatin monotherapy. This highlights an important characteristic of CDA: the worst-case scenario for any combination is that the combination performs just as well as the best-performing monotherapy at a given point in time.

The most important aspect of the tCDA model is that it represents a baseline model we can test for non-additive combinations. Non-additivity includes antagonism and synergy, the latter being more important for developing clinically successful combinations. 39% of the tested combinations can be classified as non-additivity, we reject the null model in these cases (Fig. 2A). It is critically important to understand these non-additive combinations, as they represent the combinations that perform much better or worse than expected compared to their monotherapy results. Despite an empirically (measured by root mean square distance) and visually good fit, we must be careful in jumping to the conclusion that tCDA describes the effect of the actual combination. This phenomenon can be noted in the Dabrafenib and Trametinib combination in metastatic BRAF-mutant cutanous melanoms (Supp. Fig. 5B). The tightness of the fit acts as a red herring, distracting us from the fact that it is in fact a non-additive interaction (p-value < 0.01). The non-additivity encouragingly suggests a synergistic interaction which is supported by the mechanism of action of both drugs. Trametinib is a selective inhibitor of MEK1/2

activity and Dabrafenib is a potent inhibitor of BRAF and CRAF [22, 23]. Both BRAF and MEK are in the same pathway where BRAF is upstream of MEK, so drugs that inhibit these two targets provide a classical example of synergy. There are more obviously synergistic combinations, like the 5-FU and Oxaliplatin combination for advanced pancreatic cancers for which the goodness-of-fit testing framework supports the visual fit and mechanistic understanding of the combination (Supp. Fig. 3A).

Dose-space Correlated Drug Action is a generalized additivity model at the level of cells in culture

In this section we investigate the principle of independent drug action at the level of cell cultures and in dose space. To do so we apply the same principles that were applied for tCDA but rather than asking for the survival time of a patient under each of two drugs, we will ask for the lethal dose of a cell to each of two drugs at their respective concentrations after a fixed time has elapsed since treatment. The assumption in most of what follows is that the lethal doses at which a cell dies with one or another drug are correlated. This correlation could be due to some degree of drug interaction of to confounding variables. In cell cultures we can think that specific characteristics of a cell such its number of mitochondria [24], its size, its protein content, etc, could be a confounder. We will call this application of correlated behavior of cells in dose space the Dose Correlated Drug Action, or dCDA. The dCDA model can be expressed as the following closed-form expression (see Supplement for its derivation)

$$\begin{aligned} V_{AB}(D_A, D_B, \rho) &= (1 - |\rho|)V_B(D_B)V_A(D_A) \\ &\quad + |\rho| \begin{cases} \min(V_A(D_A + g(D_B)), V_B(D_B + f(D_A))) & \text{if } 0 \leq \rho \leq 1 \\ \max(0, V_A(D_A) + V_B(D_B) - 1) & \text{if } -1 \leq \rho \leq 0 \end{cases} \end{aligned} \tag{2}$$

where $V_{AB}(D_A, D_B, \rho)$ is the dCDA predicted viability of the cell culture for the combination of drugs A and B at doses D_A and D_B when the Spearman's rank correlation of the lethal doses to A and B is ρ, $V_A(D_A)$ and $V_B(D_B)$ are the dose response curves for the cell culture treated with drugs A and be independently. The function $f(D_A)$ (resp. $g(D_B)$) is the equivalent dose of A (resp. B) that would produce the same effect in the culture viability as the dose D_B of B (resp., the dose D_A of A) did, and are discussed in the Supplement.

The dCDA framework accommodates a spectrum of models (see Supplement). For $\rho = 0$ the dCDA model reduces to the Bliss additivity model. For $\rho = 1$ the dCDA model reduces to a modified version of the HSA model that is sham compliant. Indeed the sham compliance should only be expected when $\rho = 1$, when a drug is "combined" with itself, the lethal doses of a cell to the same drug would result in a value of $\rho = 1$.

The dCDA model requires that we find the optimal parameter ρ in order to estimate the combination viability across a range of different doses. Across the 26 experimentally tested combinations discussed

below, the average Spearman's correlation estimate is 0.03 and ranges from -0.14 to 0.21 (Fig 3A). For those combinations for which the Spearman's correlation estimate is very close to zero, the dCDA model reduces to a considerable degree to the Bliss additivity model. For a given combination estimate, we can assess whether the dCDA model sufficiently describes the observed combination via a goodness of fit (GoF) test, specifically a paired two-sample t-test (Fig. 3B). 80.7% of the combinations were describable under the sham compliant dCDA model (Fig. 3A double-dagger labeled bars). Additionally, the GoF can be visually inspected by observing the extent to which the ordered pair of observed and estimated combination viabilities using dCDA lie along the identity line (Fig. 3D,F).

It is plausible, but not necesary, that the Spearman's correlation estimate in the dCDA model describes to some extent the similarity between mechanisms of action of the individual drugs that constitute the combination. If that were the case, the larger the absolute value of the Spearman's correlation estimate, the more similar the mechanisms of action. Since our Spearman's correlation estimates are on average not far from zero, we would expect that combinations describable under dCDA (GoF p-value > 0.01) have individual drugs whose mechanisms of action are effectively unrelated. For example, Tamoxifen targets the estrogen receptor in breast cancers whereas GW 8510 is a cyclin-dependent kinase inhibitor investigated with respect to colorectal cancer [25, 26]. These two mechanisms of action are not related in any intuitive manner. This lack of a relationship is represented by a Spearman's correlation estimate close to zero.

The majority of the tested combinations were assayed at 24 hours after administering the drugs, but a select few were collected in a time-series experiment. There are two sets of combinations for which time-series viability data was collected. Experiments 24, 25, and 26 correspond to Tamoxifen and Withaferin A at 12, 24, and 48 hours respectively (Supp. Fig. 17). The Spearman's correlation estimate was stable over time but the GoF p-value decreased over time (0.68, 0.15, 3.3e-5). Thus, the dose-space CDA additivity model sufficiently described the combination until the 48 hour time point. Experiments 21, 22, and 23 correspond to Tamoxifen and Mefloquine at 12, 24, 48 hours respectively (Supp. Fig. 16). The Spearman's correlation estimates varied over time (0.12, -0.05, 0.02) but the GoF p-values were consistently above the threshold for rejection of the dose-space null additivity model (p-value > 0.01). From both time-series experiments, we can observe that there are undiscovered temporal dynamics at play with respect to the action of drug combinations in cell culture.

Independent of the aforementioned six experiments, we collected data for Tamoxifen and Mefloquine at 24 hours in Experiment 3 and data for Tamoxifen and Withaferin A at 24 hours in Experiment 6. As part of the 20 non-time series trials, we didn't collect pure monotherapy viabilities but rather the viability at which the other drug in the combination was given in the lowest tested dosage. Thus, we can assess the robustness of the model to non-ideal input data by comparing Experiment 3 to Experiment 22 and Experiment 6 to Experiment 25. The Spearman's estimate differs by 0.05 between Trial 3 and 22 and in both cases the dCDA model sufficiently describes the true data (Fig. 3A). The Spearman's estimate differs by 0.01 between Trial 6 and 25 and once again, the dose-space CDA model p-value is concordant between trials (Fig. 3A). These results favorably suggest that non-ideal input data does not affect the

performance of the model which is important for analyzing the remaining trials' data collected in this fashion.

Excess over Bliss (EOB), one of the most popular performance metrics for assessing the synergistic or antagonistic behavior of a combination at particular dosages, is based on Bliss independence. In order for EOB to be a valid measure, the underlying combination must be describable under Bliss independence. Given that the dose-space CDA model collapses to Bliss independence when $\rho = 0$, we can formally test whether the combination can be sufficiently explained by Bliss independence (See Methods). The EOB condition is valid for 42.3%, 11 of the 26, of tested combinations (Fig. 3A, asterisk-labelled bars). Intuitively, the closer the optimal Spearman's correlation estimate is to zero, the more likely the EOB condition is to be valid. We can also investigate the EOB and GoF status for each tested combination (Fig. 3B). At the intersection, 38.5% of combinations are both describable under dose-space CDA and satisfy Bliss independence (Fig. 3B, blue dots).

The discussion above suggest a measure of non-additivity that we will call Excess Over CDA (EOCDA) as the difference between the optimal estimated combination viability under dCDA and the observed combination viability. EOCDA can be visualized as a heatmap where negative values (more blue) indicate an underestimation of the true viability and positive values (more red) indicate an overestimate of the true viability (Fig. 3C, E). In this paradigm, given that the CDA estimate fits sufficiently well to the observed variability (e.g. combinations in Fig. 3C, E), underestimation corresponds to possible antagonistic drug combinations and overestimation corresponds to possible synergistic drug combinations. Of particular note, for the cases presented in Fig. 3C, E and for the vast majority of cases presented in Supp. Figs. 10-17, only a few doses per combination display synergistic or antagonistic behavior as evidenced by the number of moderately to strongly colored blue/red tiles in the heatmaps.

Discussion

When considering the action of drugs in combination to treat cancer there are two main fields of study: patient populations, and cell cultures. Patient populations refers typically to humans and mice. Cell culture refers to sub-organismal experimental setups including cell lines and organoids. We present Correlated Drug Action (CDA) as a baseline model of additivity on both of these levels. Specifically, we introduce the tCDA model for patient populations and the dCDA model for cell cultures. Both models can be formally interpreted as null models in their respective problems against which we can test whether a given combination can be described as additive. Critically, a low p-value (below 0.01) implies that we can reject the null hypothesis of CDA additivity. But, we are unable to note the specific form of non-additivity displayed in these situations. However, assessing additivity or non-additivity is important in better understanding combination therapies. Additionally, the temporal and dose-space CDA models essentially reduce down to a problem of parameter estimation whose optimal value may suggest a biological implication either at the level of cell cultures or patient populations. To implement

the CDA model the PFS curves for each monotherapy is required for tCDA, and the dose response curves for the monotherapies are required for dCDA.

At the level of patient populations, CDA assumes non-interaction between the two monotherapies in a combination with respect to cell, tissue, or organ level interactions, drawing inspiration from the independence-based approach presented in recent work [7]. CDA does account for the heterogeneity in responses across a patient population. In the temporal CDA model, the fitted parameter ρ corresponds to the Spearman's correlation, ρ, between the survival times of each patient under the individual monotherapies. In order to generate paired survival times with a given Spearman's correlation, we developed two general simulation methods (See Supplement), to compare the effect of the underlying joint distribution of survival times. There are many different ways to generate bivariate joint distributions that have a given correlation ρ and known marginal distribution; but in our case it is not possible to know which of the possible distributions is the true distribution by which nature operates. Thus, we simulated two possible distributions (the coin and window swap methods) and compared them. The coin and window swap methods of simulation are concordant on simulated data but the noise inherent within real data decreases the degree of concordance between the two methods (Supp. Fig. 7). Other studies have developed simulation based temporal frameworks [7]. Given that the differences we observed between two quite different simulation models were not essential, in this paper we choose a framework that allows us to express the model using a closed-form mathematical expression. We called this model the tCDA model, and we show (See Supplement) that it is equivalent to the coin method of simulation and also covers a majority of the variance of the window swap method of simulation (Supp. Figs. 6, 7). Overall, these results suggest that although there are many ways to created survival time pairs with a given correlation, the tCDA model may captures a large portion of the variance of other reasonable methods. The tCDA model offers a computational advantage in terms of speed and efficiency. Intuitively, we can understand this by realizing that the simulated based process requires a greater number of operations than numerically optimizing a closed-form expression. Additionally, the tCDA model is easier for non-domain experts to interpret, and offers a robust hypothesis testing procedure that can assess whether a given combination is additive or non-additive.

The tCDA model can describe clinically relevant combination cancer therapies (Fig. 1,2A); the associated hypothesis testing framework allows us to identify non-additive combinations such as Nivolumab and Ipilumumab in untreated melanoma. For the combinations that were sufficiently describable under tCDA, the Spearman's correlation estimate can be interpreted in a biological context. If we estimate a moderate to high absolute value of the Spearman's correlation (Fig. 2F; Supp. Figs. 3C, 4A, 5A), it suggests that there may be a stratification of patients such that we need only to give them the drug with better survival for them. Contrarily, when the Spearman's correlation is near zero (Figs. 1, 2C, Supp. Fig. 3D), it suggests that there is no common latent variable affecting the action of both monotherapies. Under this scenario, we recognize that the combination has maximal benefit since there is no relationship between the two monotherapies' survival times. Most importantly, in these situations, we are able to leverage our uncertainty by giving the entire combination to each patient. We have no

method to predict which patient will respond better to one of the monotherapies in the combination and thus, as a hedge against our uncertainty, it may be convenient to give both therapies to each patient.

The progression-free survival (PFS) curve has a resolution equivalent to the number of patients in the trial. Since trial size defines the resolution of the survival curve, trials with very low sample sizes had poor PFS curve estimates with respect to resolution. In particular, the trial with 5-FU and Oxaliplatin combination in advanced pancreatic cancers had less than twenty patients (Supp. Table 3). Thus, the range of possible survival curves was less flexible and contributes to the inability of tCDA to explain the effect of the combination (Supp. Fig. S3A). In order to collect clinical trial data that was in the format necessary for the tCDA model, we often had to aggregate data from up to three different clinical trials (Supp. Table 3). In these scenarios, we attempted to match cohort characteristics including sex, age, ethnicity, and previous treatment history as well as possible.

There are two key priorities when evaluating any anti-cancer therapy: a) maximize tumor killing, b) minimize off-target effects and toxicity. We primarily investigated the former constraint and did not focus on the latter constraint. Both are equally important. For instance, we can assemble a combination of every anti-cancer therapy ever approved and it may do an excellent job killing the tumor but that does not matter if it is so toxic that it causes extensive damage to healthy tissue.

Transitioning from the space of patients to cell lines, the dCDA model acts at the level of cells in culture and, like the tCDA model, the dCDA model can be formulated as a closed-form expression. In this space, there have been many proposed models of additivity with Loewe and Bliss being the most used. The introduced dCDA model is inclusive because it incorporates different additivity models and concepts seamlessly. When the Spearman's correlation is zero, the model reduces down to Bliss additivity but when the Spearman's correlation is one, the model takes into consideration the Dose Equivalence principle (of which Loewe additivity is one example) (See Supplement). Furthermore, the dCDA model verifies the sham combination compliance. This defines a spectrum wherein varying the Spearman's correlation from zero to one controls whether the combination is more Bliss-like or more Loewe-like. Thus our paradigm is consistent with the Saariselkä agreement [20]. We also developed another dose-space CDA model, which we will call Bliss-HSA model, where the Spearman's correlation ranging from zero to one describes a spectrum from Bliss to the Higher Single Action model (HSA), by not assuming (as done in the dCDA model) that the effective dose of a given drug changes because of the presence of the other (See Supplement). The results of this Bliss-HSA modified dose-space CDA model were extremely consistent in terms of EOCDA, goodness of fit, and other computed statistics with the dCDA model (Supp. Table 2).

By assessing goodness-of-fit of the estimated combination viabilities under the dose-space CDA model, we can assess whether the given combination behaves in an additive manner or not (Fig. 3A). Additionally, we can apply our hypothesis testing framework to assess whether the Excess Over Bliss (EOB) performance measure of synergy and antagonism is valid in a given combination (Fig. 3A). The joint distribution of p-value statistics for the tested combinations illustrates the degree of relatedness between

the two metrics (Fig. 3B). We can notice that more cases are describable under dose-space CDA over Bliss additivity (more blue and green dots than blue and yellow dots in Fig. 3B). Thus, as a more general alternative to EOB that is valid in a wider range of conditions, we introduce Excess Over CDA (EOCDA). EOCDA is simply the difference between the estimated and observed combination viabilities under dCDA; positive values correspond to synergistic combinations and negative values correspond to antagonistic combinations (Fig. 3C,E). Interestingly, for combinations describable under dCDA, there are only a few high magnitude EOCDA values and these tiles identify the doses that are selectively synergistic.

The fitted parameter ρ in the dose-space model corresponds to the Spearman's rank correlation between the lethal doses of two different therapies for one cell across the cell culture. The Spearman's correlation provides us information regarding the relatedness between the two monotherapies including the possibility of the two drugs to target the same pathway. The reason we are able to discuss possible biological explanation from a much more general definition in the cell line case is that unlike patients, cell lines are relatively homogeneous tumor cells free of the external milieu that exists in a human body. A high positive correlation between two drugs at the dose-space level may be indicative of a shared mechanism of action between the drugs. A near zero or weak correlation implies that the mechanisms of action of both drugs are not related. A high negative correlation implies, rather interestingly, that the mechanisms of action are polar opposites of one another. In our data, all of the correlation estimates are weak implying that, from the perspective of dCDA, the drugs probably act using althernative pathways (Fig. 3A). For example, GW 8510 is a cyclin-dependent kinase inhibitor whereas Withaferin A inhibits Vimentin, an intermediate filament. These two mechanisms are not linked by any known direct relationship, and this supports the low Spearman's correlation estimate given by the dose-space CDA model (Fig. 3E, F).

Cell line experiments are often the starting ground from which researchers test and validate combinations with the eventual purpose of applying it to patients in the form of a clinical trial. There have been many discussions regarding the limitations of cell lines in representing the complexity, heterogeneity, and structure of tumors within patients [27]. This limitation is evidenced in the differences in interpretation of the Spearman's correlation estimates for the dose-space and temporal CDA model. The dose-space Spearman's correlation measures some degree of relationship between shared targets in a single clone of tumor cells. But, the Spearman's correlation at the temporal scale captures not only the mechanisms of action, but also the heterogeneity of the tumor, the immune response, and other bodily variables that are not present in a cell culture.

We can apply the tCDA framework to impute the effect of no treatment or placebo in a cohort. Placebo treatment defines a true baseline and any benefit of such a treatment (e.g. inert pills, sham surgery, or inert injections) is due to a patient's belief in the treatment [28]. The combination of placebo and an anti-cancer drug is, by definition, an instance of CDA where the Spearman's correlation is zero because there is no possible mechanism for interaction of the two therapies in a patient's body. Thus, we can estimate the placebo effect for a given trial by reversing the CDA model and assuming that the distribution form of the placebo PFS curve is a Hill curve (Supp. Fig. 18). To estimate, we

must collect data for a given anti-cancer drug and the anti-cancer drug along with the placebo (See Methods). For certain combinations (Supp. Fig. 18B), the placebo PFS curve is unrealistic as it goes to zero almost immediately. Such a placebo estimate is to be interpreted as the placebo not providing any benefit to the population of patients. The estimated placebo PFS curves would probably be more informative if we used survival and not PFS as the primary endpoint. In other combinations (Supp. Fig. 18 A,C), the placebo PFS curve estimate is plausible and highlights the relative benefit of any of the anti-cancer therapies. This foray into estimating the placebo effect is a unique use case of our null additivity temporal CDA model.

There are many possible avenues for extending this work. We can develop a network for a given cancer. Each node is a particular monotherapy and the edge width between two nodes corresponds to the Spearman's correlation estimate. This could be done for cell line or patient population. From this setup, we can attempt to learn on the graph and infer the edge strength between other pairs of nodes. This could be informative in suggesting new candidates for cell line experiments or clinical trials. As an extension of the existing work, a software package or cloud tool could be developed that allows researchers to apply the temporal and dose-space CDA models to their own datasets. Additionally, we can investigate the relationship between the dose-space and temporal CDA models by testing the same cancer combination therapy in a mice PDX model and a cell line.

The CDA framework, both for the patient population and cell line case, can be extended rather naturally to combination approaches outside the field of cancer including HIV/AIDS, antibiotic resistant infections, malaria, heart disease, and other clinical conditions. The endpoint measured would shift from PFS to another disease-relevant metric like time to response, or time to resistance.

Methods

Collection of clinical trial data

We collected data from clinical trials using the clinicaltrials.gov query tool and scraped Progression Free Survival (PFS) data from their associated results and papers. We converted graphs into tables of values through a robust online digitizer (link to digitizer). Though it would be best to have the original data, these files were not available to us. For a given combination, we aimed to find a single trial that tested the individual therapies and the combination. However, we also had to pull data across different trials to construct the necessary survival data. In these cases, we tried to maximize overlap in patient characteristics, dosages, and time schedules. We were able to match dosages and time schedules in the overwhelming majority of cases. We assembled data for 18 distinct combinations spanning 10 different cancer types and a detailed guide to the raw data can be found in Supp. Table 3. The tables of the raw PFS curves from the digitized plots that we used for our analysis can be found here.

tCDA model

Code for the Temporal CDA Model, tCDA, can be found here. Results for tCDA on the clinical trials can be found in Supp. Table 1 (Sheet 1).

Placebo

There exists a survival curve for patients when given a placebo. This placebo curve is the absolute basis of comparison. One approach to imputing the placebo curve is to treat placebo as a drug. The tCDA model assumes non-interaction, and a placebo, by definition, cannot interact with other drugs. Therefore, if we have available to us the survival curve of a given drug, and the survival curve of the given drug and placebo we can impute a possible survival curve under placebo only. Since independence is assumed, the tCDA model where rho is zero is used. Then, we parametrize the placebo curve via a Hill curve with two free parameters. We optimize the values of these parameters such that the root mean square distance between the estimated drug and placebo combination curve is minimized compared to the observed values. The range for any given Hill curve parameter is very wide to incorporate our uncertainty about where the optimal estimate may lie. We move in steps of 0.1 from 0 to 6. A p-value can be computed once an estimate for the placebo is found in the same manner that was used for estimating clinical trial data (See Supplement) using 50000 samples to construct the null distribution of KS statistics.

Code for generating placebo results can be found here. A guide to where the data comes from can be found in Supp. Table 3 and the raw tables of the scraped PFS curves can be found here. In our numbering of the trials, trials 2, 4, 8, and 18 correspond to those with placebo arms. Results of the placebo estimation can be found in Supp. Table 1 (Sheet 2).

Cell Line Data

For the time-series experiments (trials 21 - 26) we used data from Diaz et. al 2020 [29]. Supplementary File 1 from this paper was used for the Tamoxifen-Mefloquine combination in MCF7 cells at 12, 24, and 48 hour time points. Supplementary File 2 from the paper was used for the Tamoxifen-Withaferin A combination in MCF7 cells at 12, 24, and 48 hour time points. The data, as we used it, can be found here.

The following section describes the procedure for the remaining trials (1-20). The protocol and methodology we used to generate the cell line data is that of Diaz et. al 2020 [29]. The rest of this paragraph describes the method, almost exactly as presented in Diaz et. al 2020. MCF7 (ATCC HTB-22) cells were obtained from ATCC. Cells were cultured according to manufacturer's recommendations in ATCC-formulated Eagle's Minimum Essential Medium (Catalog No. 30–2003) with 10% heat-inactivated fetal

bovine serum, and 0.01 mg/ml human recombinant insulin. Growth media was changed every 3–4 days. After reaching confluence, cells were split at a ratio 1:6. Cultures were tested for mycoplasma periodically using MycoAlert (Lonza, Cat No. LT07-701) per manufacturer's instructions. To split, media was removed, cells were washed with PBS, and trypsin-EDTA mix was added for 5 min. After detachment, cells were washed with growth media, collected into 50 ml vial, spin down at 1000 RPM, suspended in fresh media and plated into 75 cm flasks. The cells were plated at 10,000 cells per well in a clear bottom black 96-well plate (Greiner Cat. No. 655090) and a white 96-well plate (Greiner Cat. No. 655083) then they were placed in an incubator. After 24 hr, the plates were removed from the incubator and treated with drugs using the HP D300 Digital Dispenser. Cells were then treated with the appropriate drug compounds depending on the trial under consideration. For example, if we were testing the Emetine and Imatinib combination within MCF7 cells, we treated the cells in each well with a distinct dose of Emetine and Imatinib such that we could recover viability data for the specified concentrations as shown in Fig. 3C. After the targeted drug treatment time, 24 hours, 100 µL of Cell-Titer-Glo (Promega Corp.) was added to the wells in the white 96-well plate and shaken at 500 rpm for 5 min. The plate was then read by the Perkin Elmer Envision 2104 using an enhanced luminescence protocol to count the number of raw luminescent units per well. For the black clear bottom 96-well plates, the plate was spun at 300 g for 5 min and all the cell media was removed. Methanol was then added at 200 µL per well and let sit at room temperature for 15 min. The methanol was removed from the wells and 200 µL of PBS with Hoechst 33342 nucleic acid stain at a final concentration of 1 µG/mL was then added to the wells. The plates were then imaged with the GE Healthcare IN Cell Analyzer 2000 that is equipped with a CCD camera. The IN Cell Analyzer software was used to count the number of cells detected in each well to calculate the viability.

We tested 20 unique combinations with 13 unique drugs. For each combination, we tested different concentrations of both drugs and measured the percent of cells that were inhibited at 24 hours. Thus, we have a matrix where the (i,j)th entry corresponds to the percent of cells inhibited under a given concentration of both drugs. We did not have data for the inhibition under just one drug (i.e. other drug had a concentration of 0 μM) for trials 1-20. The data for an individual drug's inhibition was approximated by considering the case when the other drug was measured at its lowest concentration. The raw viability dose matrices for trials 1-20 can be found here.

dCDA model

Code for the dose-space CDA model, dCDA, where $\rho = 1$ corresponds to HSA can be found here. Results from this model can be found in Supplemental Table 2 (Sheet 2).

Code for the sham compliant dCDA model can be found here. Results from this model can be found in Supplemental Table 2 (Sheet 1). Note that the results presented in this paper are that of the sham compliant dCDA model. Both models, however, are highly concordant in their results.

Authors' contributions

GS conceived of the project. All co-authors developed the methodology for the project. ASA and SCK developed computational approaches and performed bioinformatics analysis with guidance from GS and MEA. ASA drafted the manuscript and all authors participated in its editing and revision.

Competing interests

The authors declare that they have no competing interests.

Availability of Data and Materials

All data and code can be found at github.com/aditharun/correlated-drug-action.

References

[1] Ann E. Kurth et al. Combination HIV Prevention: Significance, Challenges, and Opportunities. *Curr HIV/AIDS Rep.*, 8(1):450–470, 2011.

[2] Emil Frei III et al. Studies of Sequential and Combination Antimetabolite Therapy in Acute Leukemia: 6-Mercaptopurine and Methotrexate. *Blood*, 18(4):431–454, 1961.

[3] Emil J. Freireich et al. The Effect of 6-Mercaptopurine on the Duration of Steroid-induced Remissions in Acute Leukemia: A Model for Evaluation of Other Potentially Useful Therapy. *Blood*, 21(6):699–716, 1963.

[4] Reza Bayat Mokhtari, Tina S Homayouni, Narges Baluch, Evgeniya Morgatskaya, Sushil Kumar, Bikul Das, and Herman Yeger. Combination therapy in combating cancer. *Oncotarget*, 8(23):38022, 2017.

[5] Emil Frei III et al. The Effectiveness of Combinations of Antileukemic Agents in Inducing and Maintaining Remission in Children with Acute Leukemia. *Blood*, 26(5):643–656, 1965.

[6] Jonathan B. Fitzgerald et al. Systems Biology and Combination Therapy in the Quest for Clinical Efficacy. *Nature Chemical Biology*, 2(9):450–470, 2006.

[7] Adam C. Palmer and Peter K. Sorger. Combination Cancer Therapy Can Confer Benefit via Patient-to-Patient Variability without Drug Additivity or Synergy. *Cell*, 171(7):1678–1691, 2017.

[8] William R Greco. The search for synergy: a critical review from a response surface perspective. *Pharmacol Rev*, 47:331–385, 1995.

[9] CT Meyer, DJ Wooten, BB Paudel, J Bauer, KN Hardeman, D Westover, et al. Quantifying drug combination synergy along potency and efficacy axes. cell syst. 2019; 8: 97–108. e16.

[10] Mark Sinzger, Jakob Vanhoefer, Carolin Loos, and Jan Hasenauer. Comparison of null models for combination drug therapy reveals hand model as biochemically most plausible. *Scientific reports*, 9(1):1–15, 2019.

[11] Cong Chen, Fang Liu, Yixin Ren, Leah Suttner, Zhiping Sun, Yue Shentu, and Emmett V Schmidt. Independent drug action and its statistical implications for development of combination therapies. *Contemporary Clinical Trials*, 98:106126, 2020.

[12] C.I. Bliss. The Calculation of Microbial Assays. *Bacteriol Rev.*, 20(4):243–258, 1956.

[13] S Loewe. The Problem of Synergism and Antagonism of Combined Drugs. *Arzneimittelforschung*, 3(6):285–290, 1953.

[14] Julie Foucquier and Mickael Guedj. Analysis of drug combinations: current methodological landscape. *Pharmacology research & perspectives*, 3(3):e00149, 2015.

[15] M.C. Berenbaum. What is Synergy? *Pharmacol Rev.*, 41(3):422, 1989.

[16] Vincent T. DeVita Jr and Edward Chu. A History of Cancer Chemotherapy. *Cancer Res.*, 68(21):8643–8653, 2008.

[17] Justin R. Pritchard et al. Understanding Resistance to Combination Chemotherapy. *Drug Resist Updat.*, 15(0):249–257, 2012.

[18] Vincent T. DeVita Jr and Steven A. Rosenberg. Two Hundred Years of Cancer Research. *N Engl J Med.*, 266(23):2207–2214, 2012.

[19] George P. Canellos et al. Chemotherapy of Advanced Hodgkin's Disease with MOPP, ABVD, or MOPP Alternating with ABVD. *N Engl J Med*, 327(21):1478–1484, 1992.

[20] Jing Tang et al. What is Synergy? The Saariselkä agreement revisted. *Front Pharmacol.*, 6(181), 2015.

[21] D.J Slamon et al. Use of chemotherapy plus a monoclonal antibody against HER2 for metastatic breast cancer that overexpresses HER2. *N Engl J Med.*, 344(11):783–792, 2001.

[22] Iwona Lugowska et al. Trametinib: a MEK Inhibitor for Management of Metastatic Melanoma. *Onco. Targets Ther.*, 8(1):2251–2259, 2015.

[23] Samantha Bowyer et al. Dabrafenib and Its Use in the Treatment of Metastatic Melanoma. *Melanoma Manag.*, 2(3):199–208, 2015.

[24] Luís C Santos, Robert Vogel, Jerry E Chipuk, Marc R Birtwistle, Gustavo Stolovitzky, and Pablo Meyer. Mitochondrial origins of fractional control in regulated cell death. *Nature communications*, 10(1):1–10, 2019.

[25] M.M. Gottardis and V.C. Jordan. Antitumor Actions of Keoxifene and Tamoxifen in the N-nitrosomethylurea-induced Rat Mammary Carcinoma Model. *Cancer Res.*, 47(15):4020–4024, 1987.

[26] Y-Y Hsieh et al. Repositioning of a Cyclin-Dependent Kinase Inhibitor GW8510 as a Ribonucleotide Reductase M2 Inhibitor to Treat Human Colorectal Cancer. *Cell Death Discov.*, 2, 2016.

[27] Jennifer L. Wilding and Walter F. Bodmer. Cancer Cell Lines for Drug Discovery and Development. *Cancer Res.*, 74(9):2377–2384, 2014.

[28] Tor D. Wager and Lauren Y. Atlas. The Neuroscience of Placebo Effects: Connecting Context, Learning, and Health. *Nat Rev Neurosci*, 16(7):403–418, 2015.

[29] J.E. Diaz et al. The Transcriptomic Response of Cells to a Drug Combination is More Than The Sum of the Responses to the Monotherapies. *eLife*, 9, 2020.

Figures

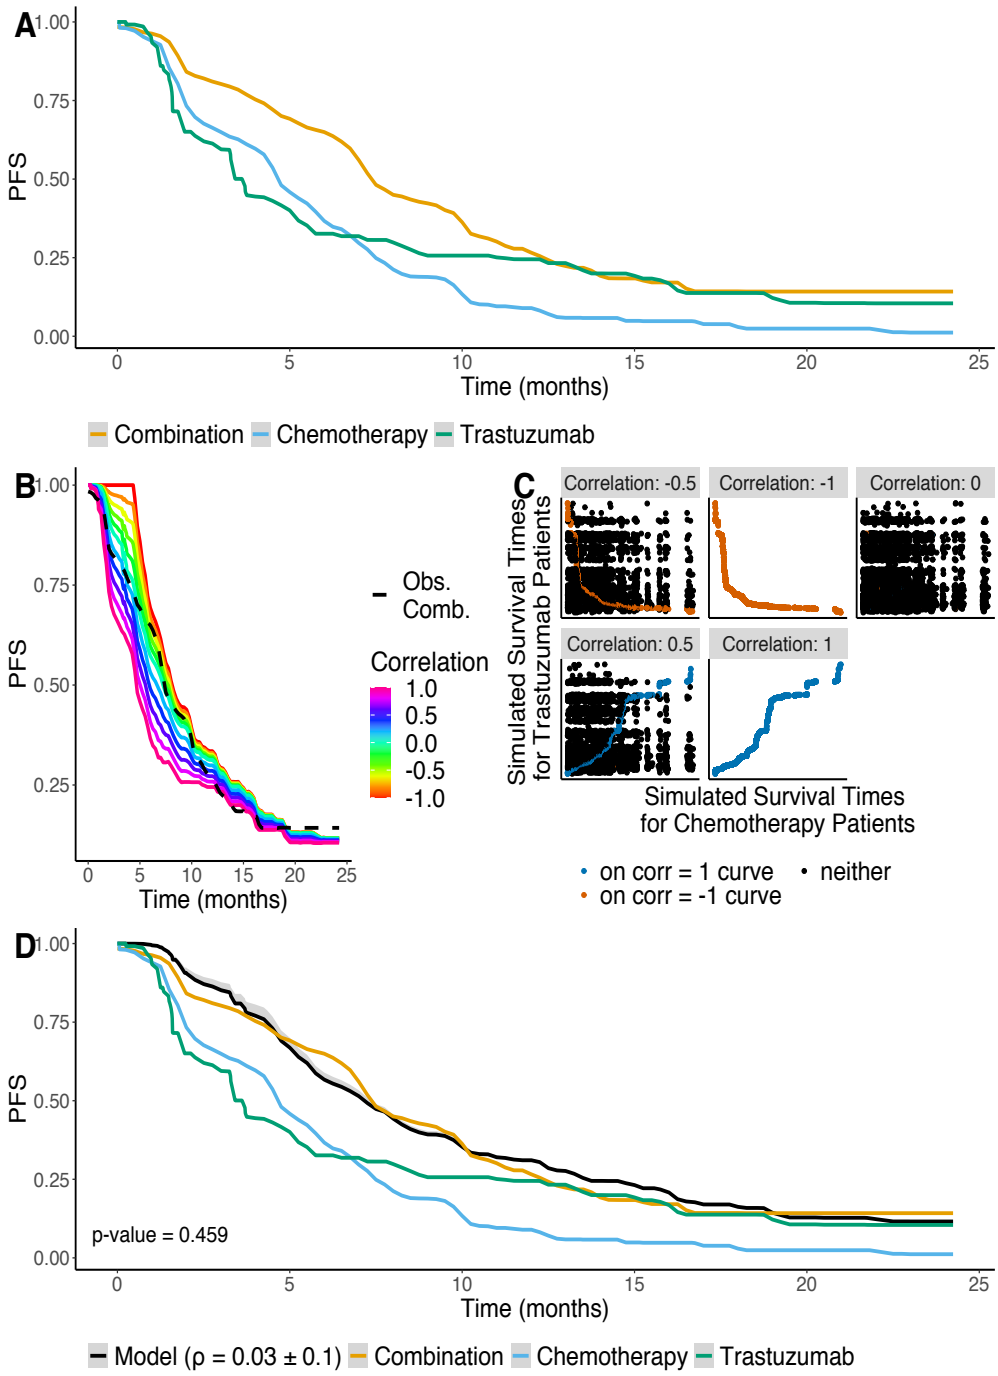

Figure 1: Correlated Drug Action explains the benefit of Trastuzumab and Chemotherapy in Metastatic HER-2 Overexpressing Breast Cancer. A) Progression Free Survival (PFS) as a function of time for either individual therapy option (Trastuzumab - green, Chemotherapy - blue) and their combination (yellow) in patients with HER-2 overexpressing breast cancer. **B)** Range of possible survival curves for the combination under tCDA. The observed combination (black) falls within the lines of the predicted field. **C)** For a given Spearman's correlation, each point represents a pair of possible PFS times associated with each simulated patient, and the parallel maximum of this paired vector of PFS times defines the combination PFS curve shown in **B**. The absolute value of the Spearman's correlation corresponds to the fraction of points that lie on the correlation -1 or 1 curves (Supp. Figure 8). **D)** Estimate of the combination under the tCDA model (black) and in grey are the PFS curves for the 95% confidence interval of ρ_s (0.03 $\pm$ 0.1, p-value $=$ 0.459). This suggests that the tCDA additivity model sufficiently describes the effect of the combination.

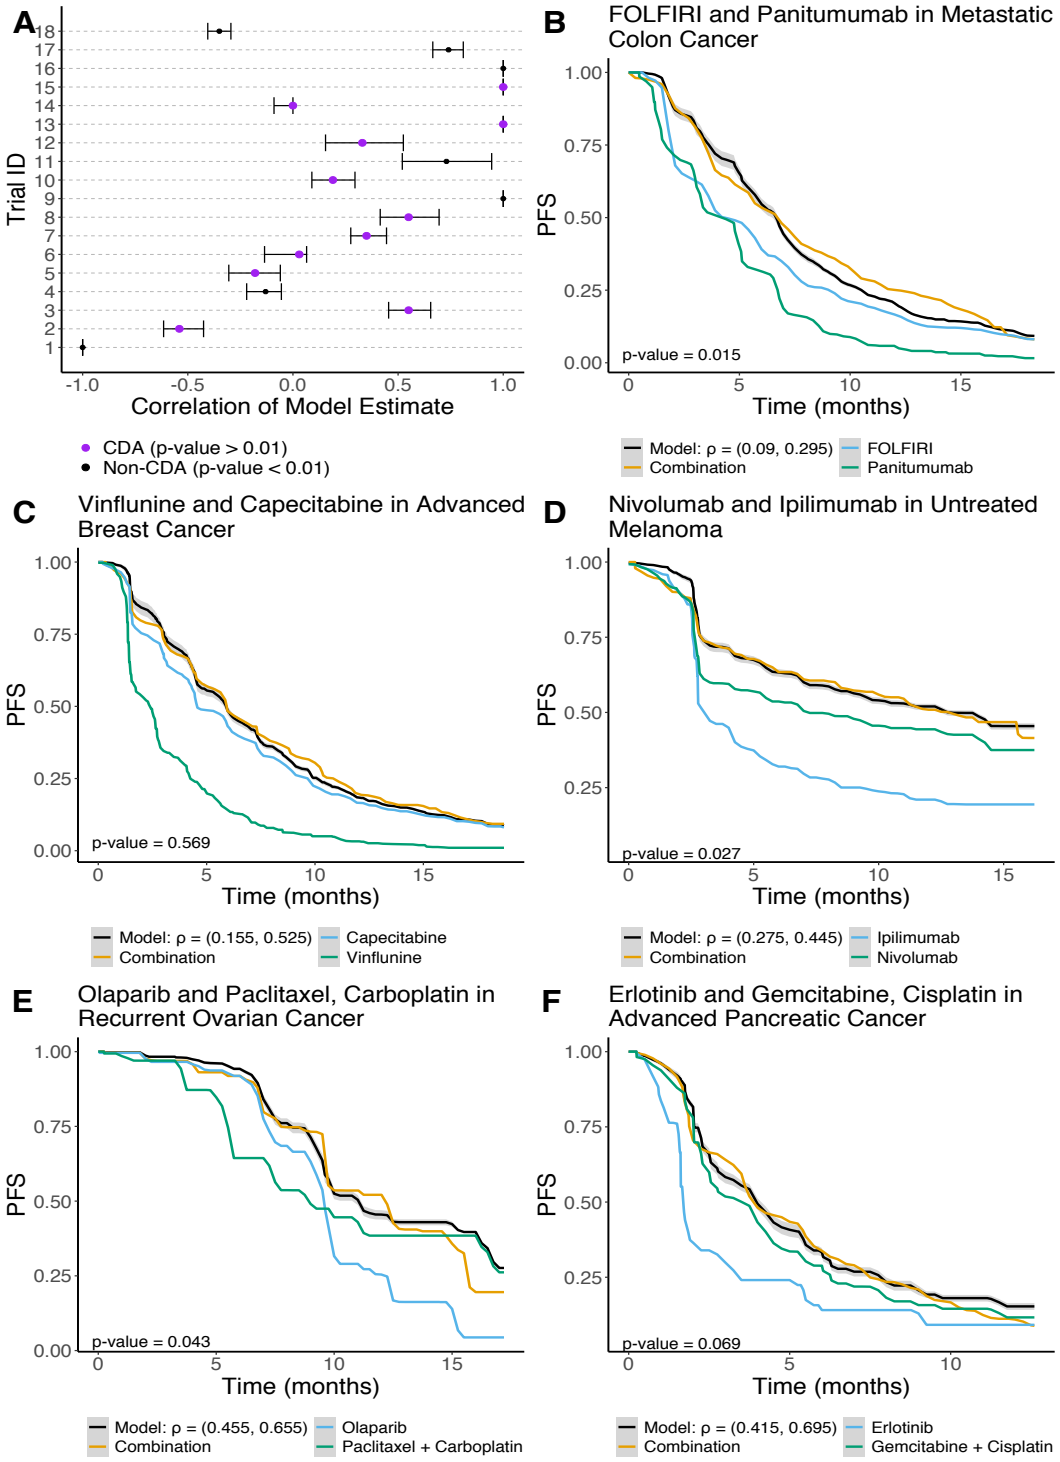

Figure 2: Temporal Correlated Drug Action (tCDA) model results. A) Optimal Spearman's correlation and associated 95% confidence interval. Purple represents trials where the effect of the combination can be explained sufficiently well under the tCDA model. The trial ID's reference the tested clinical trials, and a lookup table can be found in Supp. Table 1. **B)** tCDA model estimate for the combination of FOLFIRI and Panitumumab in metastatic colon cancer. **C)** tCDA model estimate for the combination of Vinflunine and Capecitabine in advanced breast cancer. **D)** tCDA model estimate for the combination of Nivolumab and Ipilumumab in previously untreated melanoma. **E)** tCDA model estimate for the combination of Olaparib and Paclitaxel, Carboplatin in recurrent ovarian cancer. **F)** tCDA model estimate for the combination of Erlotinib and Gemcitabine, Cisplatin in advanced pancreatic cancer.

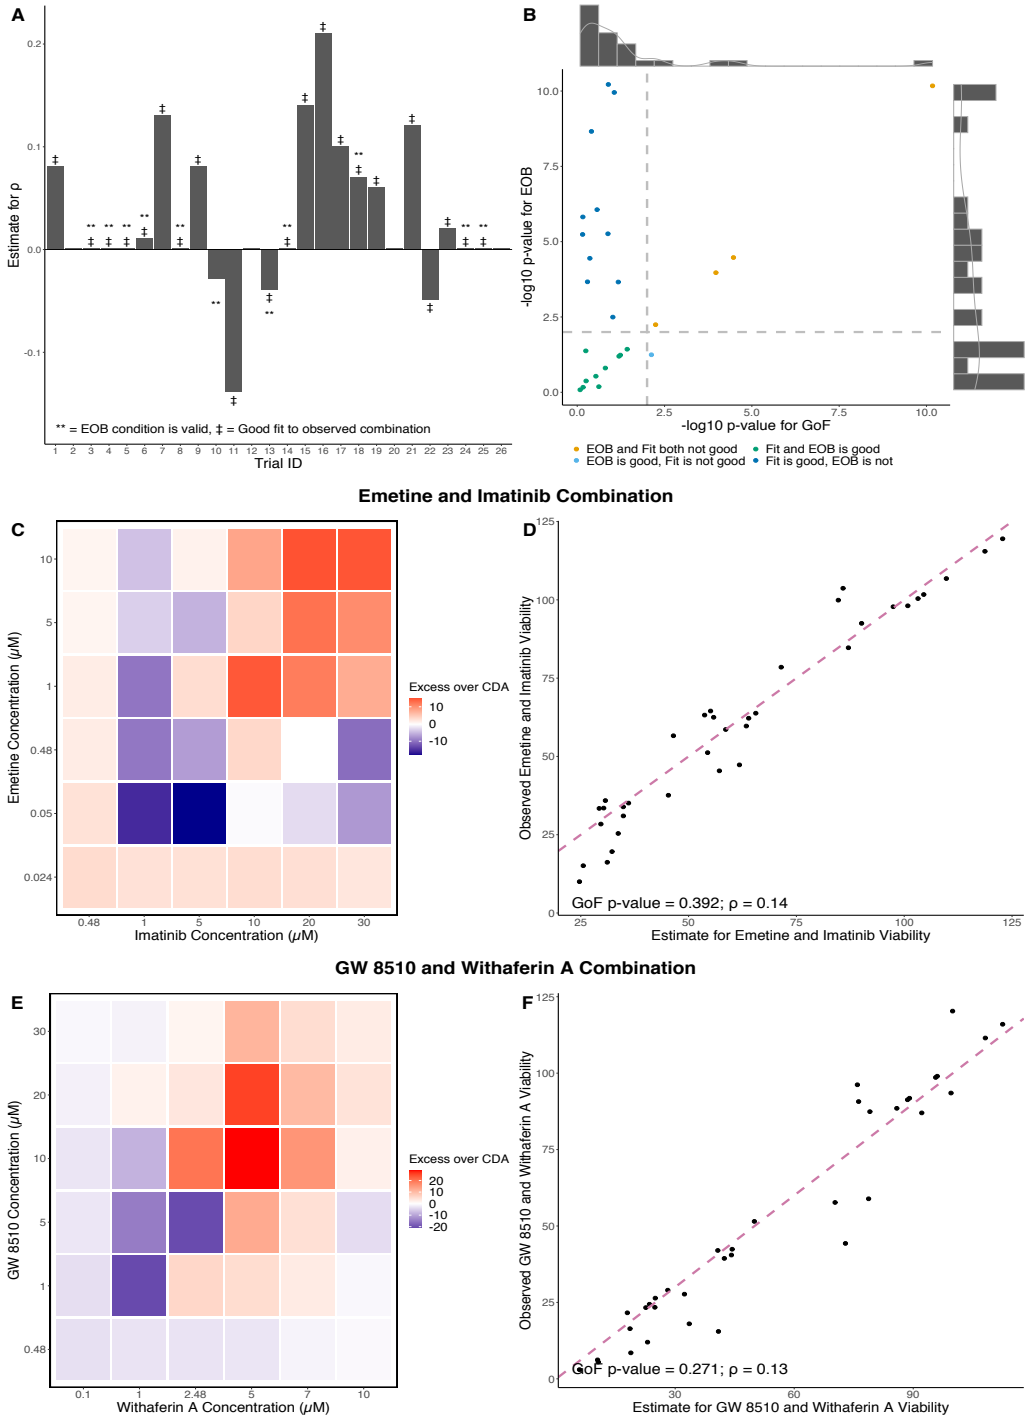

Figure 3: Sham compliant dose-space Correlated Drug Action (dCDA) model results. A) Optimal Spearman's correlation estimate for each *in vitro* drug combination tested on MCF7 cells under the dCDA model. Double-dagger indicates a good fit (p-value > 0.01), and star indicates that Excess Over Bliss (EOB) is a valid metric (p-value > 0.01). The trial ID's correspond to a given drug combination (Supp. Table 2). **B)** Joint distribution of p-value statistics for Goodness of Fit and EOB conditions (Dashed lines = p-value of 0.01). **C, D)** Emetine and Imatinib combination results. **C)** Heatmap of the EOCDA, predicted minus observed combination viabilities, metric. Positive values (more red) suggest synergy whereas negative value suggest antagonism (more blue). **D)** Comparison of estimated and observed viabilities of Emetine and Imatinib in reference to the identity line (p-value = 0.392). **E, F)** GW 8510 and Withaferin A combination results. **E)** Heatmap of EOCDA. **F)** Comparison of estimated and observed viabilities (p-value = 0.27).